

UNIVERSITY OF RIJEKA
FACULTY OF HUMANITIES AND SOCIAL SCIENCES
DEPARTMENT OF PHILOSOPHY

Marko-Luka Zubčić

**CERTAIN FUNDAMENTALS
OF THE DESIGN AND ASSESSMENT
OF SOCIAL EPISTEMIC SYSTEMS:
AGENTS SUBOPTIMALISM,
INSTRUMENTAL PLURALISM
AND UNIVERSAL INCLUSION**

DOCTORAL THESIS

Rijeka, 2020.

UNIVERSITY OF RIJEKA
FACULTY OF HUMANITIES AND SOCIAL SCIENCES
DEPARTMENT OF PHILOSOPHY

Marko-Luka Zubčić

**CERTAIN FUNDAMENTALS
OF THE DESIGN AND ASSESSMENT
OF SOCIAL EPISTEMIC SYSTEMS:
AGENTS SUBOPTIMALISM,
INSTRUMENTAL PLURALISM
AND UNIVERSAL INCLUSION**

DOCTORAL THESIS

Mentor: prof. dr. sc. Snježana Prijić-Samažrija

Rijeka, 2020.

Mentor rada: prof. dr. sc. Snježana Prijic Samaržija

Doktorski rad obranjen je dana _____ u _____,
pred povjerenstvom u sastavu:

1. _____
2. _____
3. _____

Abstract

The thesis examines foundational arguments of design and assessment of social epistemic systems, an area of epistemology which studies epistemic properties of social and institutional arrangements. First chapter presents the historical overview of institutional epistemology, focusing on pragmatism, experimentalism and democracy in the work of John Dewey, ignorance, norms, pluralism and market in the work of Friedrich Hayek, and the contemporary use of simulations in epistemological research. Second chapter condenses the advances in the discipline: (i) comparative standard for the assessment of social epistemic systems is defined; (ii) baseline conditions of the epistemic life of the population are defined; (iii) instrumentalist arguments for normative pluralism and universal inclusion as the primary features of epistemically superior social epistemic systems are derived from the two constraints of design. Third chapter opens a discussion on the index of epistemic progress by which the features would be assessed.

Keywords

epistemology / institutions / pluralism / inclusion / ignorance

Prošireni sažetak

Ovaj se rad bavi temeljnim argumentima o dizajnu i procjeni društvenih epistemičkih sustava, dijelu epistemologije koji istražuje epistemička svojstva društvenih i institucijskih uređenja.

Prvo poglavlje sadrži povijesni pregled institucijske epistemologije, s fokusom na teme pragmatizma, eksperimentalizma i demokracije kod Johna Deweya, teme neznanja, normi, pluralizma i tržišta kod Friedricha Hayeka, te suvremena istraživanja u disciplini obilježena upotrebom simulacija kao metodološkom inovacijom, kroz rad Scotta Pagea, Michaela Weisberga, Ryana Muldoona, i Kevina Zollmana.

Drugo poglavlje sažima napretke u disciplini pregledane u prvom poglavlju i nudi novu formulaciju temeljnih argumenata. Prvo, definira se komparativni standard za procjenu društvenih epistemičkih sustava – nakon pregleda, kritike i odbacivanja „istine“ i „razvoja agenta“, brani se „sposobnost sustava da revidira lažna vjerovanja“ kao komparativni standard. Drugo, definiraju se osnovni uvjeti društvenog epistemičkog sustava kao ograničenja na potencijalni dizajn – agenti su epistemički suboptimalni, vođeni normama, te njihov broj je konačan ali nepoznat. Treće, iz navedena dva ograničenja u dizajnu (komparativnog standarda i osnovnih uvjeta), izvode se instrumentalistički argumenti za normativni pluralizam i univerzalnu inkluziju kao primarna svojstva epistemički superiornog društvenog epistemičkog sustava.

Treće poglavlje otvara raspravu o indeksu epistemičkog razvoja kojime bi se navedena svojstva procjenjivala.

Ključne riječi

epistemologija / institucije / pluralizam / inkluzija / neznanje

Contents

INTRODUCTION	1
The Plan of the Thesis	2
Negative Approach	4
Preliminary Notes on Terminology	5
Institutional Epistemology and Governance: The Infrastructural View	8
1. THE PLURALIST PROJECT IN INSTITUTIONAL EPISTEMOLOGY	13
1. 1. Pragmatism, Experimentalism and Democracy: Deweyan Institutional Epistemology	14
1. 1. 1. Pragmatism	14
1. 1. 2. Experimentalism	20
1. 1. 3. Democracy	22
1. 1. 3. 1. Polycentric Governance	25
1. 1. 4. Votes and Talk	28
1. 1. 5. Conclusion	33
1. 2. Ignorance, Norms and Instrumental Pluralism: Hayekian Institutional Epistemology	34
1. 2. 1. Introduction	34
1. 2. 2. Ignorance	38
1. 2. 3. Norms	42
1. 2. 4. Instrumental Normative Pluralism	46
1. 2. 5. Prices	49
1. 2. 6. Markets and Social Epistemic Exclusion	52
1. 2. 7. Conclusion	58
1. 3. Division of Epistemic Labour and Diversity Trumps Ability: The Lessons From “Simulators”	60
1. 3. 1. Weisberg and Muldoon, “Epistemic Landscapes and the Division of Cognitive Labour”	61
1. 3. 2. Zollman, “The Epistemic Benefit of Transient Diversity”	63
1. 3. 3. Page, <i>The Difference</i>	65
1. 3. 4. Conclusion	70

2. MINIMAL PRINCIPLES OF THE DESIGN AND ASSESSMENT OF SOCIAL EPISTEMIC SYSTEM	71
2. 1. Comparative Standard in Institutional Epistemology	73
2. 1. 1. Introduction	73
2. 1. 2. The Assessment of Comparative Standards: Agent Development and Attainment of Truth	75
2. 1. 2. 1. Agent Development as the Comparative Standard	75
2. 1. 2. 1. 1. Epistemic Capabilities as Agent Development	75
2. 1. 2. 1. 2. Two Objections from the Division of Epistemic Labour	78
2. 1. 2. 2. Attainment of Truth as the Comparative Standard	79
2. 1. 2. 2. 1. Strong Political and Weak Epistemological Objection to Epistocracy	81
2. 1. 2. 2. 2. Epistemological Objection to Attainment of Truth as Comparative Standard	83
2. 1. 2. 3. Revision of a Suboptimal Epistemic State as the Comparative Standard	85
2. 1. 2. 4. Conclusion	87
2. 2. Baseline Conditions of Social Epistemic Systems: The Regulative Account	88
2. 2. 1. Epistemic agents are epistemically suboptimal	88
2. 2. 1. 1. Conservation	92
2. 2. 2. Epistemic agents are normative	94
2. 2. 3. The number of epistemic agents in the population is finite and unknown	96
2. 2. 4. Conclusion	98
2. 3. Minimal Principles	99
2. 3. 1. Redundant Normative Pluralism	100
2. 3. 1. 1. Knowledge is Minimally Conditioned on Redundant Normative Pluralism: Justification and Reliability	102
2. 3. 1. 2. Interaction	105
2. 3. 1. 3. Transience	107
2. 3. 1. 5. Conclusion	108
2. 3. 2. Universal Inclusion	108
2. 3. 2. 1. Trivial Clause	110
2. 3. 2. 2. Output Value Unpredictability Thesis	110
2. 3. 2. 3. Agent Scarcity Thesis	111
2. 3. 2. 4. Conclusion	115
2. 3. 3. Joker Objection to Minimal Principles: Particular agents may reduce the quality of the epistemic output	115
2. 3. 3. 1. The Enemies of Knowledge	116
2. 3. 3. 2. Relevant Contributors	117
2. 3. 3. 3. Unpredictable Uncertainty-reducers	117
2. 4. Conclusion	120

3. PRELIMINARY NOTES ON APPLIED INSTITUTIONAL EPISTEMOLOGY: EPISTEMIC INFRASTRUCTURE INDEX AND SUPPLEMENTARY ASSESSMENT	122
3. 1. Introduction	122
3. 2. Rationale for the Epistemic Infrastructure Index	123
3. 3. Epistemic Infrastructure Index	125
3. 4. Supplementary Assessment: Index Results Analysis and Contingent Distortions Review	131
3. 5. Conclusion	132
CONCLUSION	134
REFERENCES	136

“You gentlemen who think you have a mission
To purge us of the seven deadly sins
Should first sort out the basic food position
Then start your preaching, that's where it begins.”

Bertolt Brecht

INTRODUCTION

One of the most important discoveries in social sciences at the turn of the millenium is that collective epistemic virtues are not reducible to individual epistemic virtues. Mayo-Wilson et al (2011) call this the Independence Thesis. Groups of worse and redundant but diverse investigators epistemically outperform groups of best investigators in solving complex problems (Page 2008). This discovery underlies the contemporary development of institutional epistemology (IE), a philosophical discipline committed to interdisciplinary research into the epistemic properties of social, institutional, systems (Anderson 2006; Goldman 2011 also refers to it as “systems-oriented social epistemology”). The founding principle of IE is then that the institutional arrangement which gives rise to a redundantly pluralist populations is more conducive to knowledge then the one which would rely exclusively on those members of the population that are more likely to attain knowledge or the one which would nurture specific individual epistemic virtues in the population.

The aim of this thesis is to 1) reconstruct the discussion which has led to the establishment of IE, 2) offer a novel restatement of its foundational argument, and 3) outline a desiderata of governance reflective of this discovery. First, the development of “the pluralist project” in IE will be investigated; second, the “Minimal Principles of the design and assessment of social epistemic systems” (namely, redundant pluralism and universal inclusion) will be derived from the comparative standard (the criteria according to which systems are assessed) and baseline conditions (the general constraints on the system design) in IE; and third, the assessment of minimal “epistemic infrastructure” conducive to Minimal Principles will be presented as a research agenda and a tool for proper governance of knowledge.

In this Introduction I will firstly present the plan of the thesis. I will then address the negative approach in IE and make certain preliminary terminological clarifications. Lastly, I will briefly discuss the most urgent implications of the present account for the purposes of real-world governance.

The Plan of the Thesis

Chapter 1 will focus on the historical development of the pluralist project as foundational of IE through works of John Dewey and pragmatists in general, Friedrich August Hayek, and contemporary authors working with simulations and agent-based modelling - Michael Weisberg, Ryan Muldoon, Scott Page and Kevin Zollman.

Firstly, I will discuss pragmatism, and how its philosophical image of knowledge as constituted by the ability to experience error in judgement informs Dewey's experimentalism in IE. I will show, following Elizabeth Anderson, that Dewey found democracy to be the institutional arrangement most conducive to knowledge due to its experimentalist properties, and moreover that Elinor Ostrom's polycentric governance has the potential to mitigate a number of epistemic flaws of the institutional order exhaustively described as "democracy". Lastly, I will discuss the epistemic strengths and weaknesses of votes and talk, two mechanisms for harvesting collective intelligence representative of democracy.

Secondly, I will discuss two Hayek's foundational insights in IE - that epistemic agents are necessarily epistemically suboptimal (ignorant and cognitively limited), and that they are led in the search for knowledge by norms. I will show these assumptions, above and beyond decentralization, instruct the design of social epistemic systems towards redundant normative pluralism. Furthermore, I will discuss prices as mechanisms for harvesting collective intelligence, and lastly, markets, an institutional arrangement which I will argue fails to be conducive to knowledge because it violates redundant normative pluralism by lacking the institutional protection of redundant investigators.

Thirdly, I will discuss three classical contemporary works in IE which make use of simulations to investigate institutional and collective epistemic virtues - Weisberg and Muldoon's discovery that the presence of agents which explore unknown patches of "epistemic landscape" severely increases the capacities of groups to perform optimally, Page's (and Hong's) development of understanding that groups of diverse

investigators outperform groups of most able investigators, and Zollman's finding that conservation of suboptimal norms by diverse groups makes them more likely to collectively converge on an optimal one. These works bring necessary nuance to investigations in IE, and provide profound insights into those structures of pluralism which render it more likely to attain and produce knowledge. Some pluralisms outperform other pluralisms. In a large population faced with complex and, particularly, wicked problems (the defining feature of which is indeterminable "proximity" to optimal solution; Rittel and Webber 1973), however, any pluralism will outperform centralization.

Chapter 2 will provide a rephrasing of the central foundational argument of IE - that pluralism and inclusion are epistemically instrumental - through three steps.

Firstly, I will claim that the comparative standard according to which we may identify an epistemically better performing population is how able is it to revise a suboptimal strategy. I will argue this by presenting objections to agent development (in terms of individual virtues) and the attainment of truth as comparative standards.

Secondly, I will argue that the baseline conditions of any population of agents is that they are individually epistemically suboptimal, driven by norms and that there is a finite but unknown number of them.

I will then claim that given these two constraints in design, redundant normative pluralism and universal system-level epistemic inclusion follow. Redundant normative pluralism may be defended by Hedge Thesis, which claims that if we are suboptimal and normative, we are more likely to revise a suboptimal strategy if we follow different norms. I will furthermore argue that pluralism is minimally required for a population, or an agent, to satisfy the justification criteria of knowledge, and specify that epistemically instrumental pluralism must be interactive, and feature local transience while globally sustained (and thus, produce social epistemic inequalities while protecting from social epistemic exclusion). Universal inclusion may be defended trivially - that there is simply more of us to follow different norms if there is more of us and we are not punished for following it. It may also be defended by claiming that given that any assessors are themselves suboptimal, their judgement

of a suboptimal strategy may be wrong. And, perhaps most importantly, by claiming that since there is never enough agents to guarantee optimal epistemic outcome, each agent should be in the best position to achieve knowledge. Lastly, I will offer a response to the objection that there may be epistemic agents and communities which lower the quality of collective epistemic output.

Chapter 3 will discuss the possible application of institutional epistemology in the light of the pluralist program. I will focus primarily on an index of assessment of epistemic infrastructure in terms of universal access to sustenance, epistemic resources and possibility of Epistemic Contribution, and thus system-level epistemic inclusion. I will then address two additional steps in the assessment - first, the analysis of its results, which must be focused on the institutional inability of the found population to provide epistemic infrastructure; and second, the inquiry into the properties of populations which may be recognized as particularly epistemically distortive, particularly by the use of simulations and agent-based modelling.

Negative Approach

In her text “Epistemic Contribution as a Central Human Capability”, Miranda Fricker (2015) argues a certain *negative approach* to the design of social and institutional rules in the search for knowledge could be advised. This regulative institutional design would focus on “threats” that the social epistemic system “must stave off” (Fricker 2015). Even if we were designing an ideal social epistemic system, we would do better if we were to build it from a better “immune system” up (Fricker 2015)

I believe it is a baseline condition which constrains any design of social epistemic system that epistemic agents are necessarily suboptimal, and I will discuss it in Chapter 2, as well as in Chapter 1 in the context of Hayek’s foundational work in IE. Moreover, I will argue that the comparative standard, a criteria according to which IE

should assess social epistemic systems, should be systems' ability to revise suboptimal epistemic strategies, and thus learn. I find these understandings reflective of such a negative, regulative outlook. The system governing a population of epistemically suboptimal agents should focus primarily on the threat of it getting stuck at pursuing a wrong path in the search for knowledge.

Negative approach must however result in a positive instruction for governance, an institutional *order*, which prescribes a system-level rule translatable into a set of policies. While John Dewey believed democracy is most conducive to best epistemic outcomes and Hayek found markets to be the most promising candidate (as I will discuss in Chapter 1), I will argue that constitutional liberal democracy with free markets is necessary but insufficient for such ends. I will argue that the social epistemic system must also guarantee the protection of redundant investigators - and thus provide a specific egalitarian minimum to all, regardless of their individual performance, which I will refer to as "minimal epistemic infrastructure". I will use the term "infrastructure" in particular to connote the systemic epistemic benefit of such an institutional guarantee. It is in our collective and individual epistemic *interest* that we all have access to sustenance, epistemic resources and the possibility of Epistemic Contribution.

I will show that the Minimal Principles of the design of social epistemic systems are redundant pluralism and universal inclusion, and the failure to protect redundant investigators, which I will refer to as "social epistemic exclusion", is thus among the key threats the system must stave off in order to be more likely to attain knowledge.

Preliminary Notes on Terminology

I will here offer a rough overview of the most generally used terms in the present account of IE. More detailed analyses of certain terms will be provided at the appropriate points in the discussion.

In the minimal social ontology which I will use, the “institutional arrangement” dictates the rules that all “normative communities” of “epistemic agents”, exemplified by individuals, must follow. When it dictates rules for the search for knowledge, this institutional arrangement will be called a “social epistemic system”.

“Normative communities” here designate all groups of agents which share norms. Various normative communities may be solving the same problem and different problems. Each agent may belong to various normative communities - unless, of course, the membership in one is exclusionary of other. Epistemic agents will make “normative commitments”, epistemic actions (presumably reflective of beliefs) for which they are responsible and with which, as potentially optimal claims, local peaks in the epistemic landscape, they enter a “contest”, the adversarial social game of giving and asking for reasons (Fossen 2014).

I will not in detail discuss particular normative communities nor address the epistemic traits which may be beneficial to them. My focus will be exclusively on the rules which govern them all, and more specifically, on *minimal* epistemically beneficial rules which govern them all.

“Normative pluralism” entails a population consisting of normative communities which differ in local peaks - both in terms of differing in norms but sharing the problem which defines the epistemic task and in terms of differing in norms and thus attending to different problems. The difference in norms in numerous cases may also entail the difference in ascribing value to particular epistemic outcomes and to particular overall epistemic, aesthetic, moral and other goals, and thus *evaluative* diversity (D’Agostino 2009).

Epistemic content may be harvested from the population through various mechanisms that in turn allow for the signalling of that epistemic content “across” diverse normative communities. I will discuss “talk, votes and prices” in Chapters 1 and 2 as classical examples of such “intelligence harvest mechanisms” in IE (Anderson 2006).

Furthermore, I will in Chapter 2 argue that it is a baseline condition that agents are “epistemically suboptimal” because a) they don’t have access to all relevant evidence, b) they have limited and lacking conceptual resources, c) they make inferential mistakes and errors (systematic mistakes), d) they tend to conserve suboptimal strategies in the search for knowledge, and e) they cannot predict the future.

In this minimal ontology then every “Epistemic Contribution”, understood, following Fricker, as the “the exercise of (...) social epistemic capability on the part of the individual to contribute to the pool of shared epistemic materials – materials for knowledge, understanding, and very often for practical deliberation” (Fricker 2015, 76), is an “epistemic bet” (Muldoon and Weisberg 2011). Some bets may be better, solutions to problems may be found, and certain disagreements may be resolved. There may be “no-bet” areas - at least until the betting is on. And even where the betting is on, certain players *may* simply be excluded from some areas.

I will argue they should not however be excluded from *all*. I will argue that from the system point-of-view it is of *benefit* to the overall quality of the epistemic outcome (and thus in the collective interest of the population and individual interest of all agents) if the redundant investigators were to “remain in the game”.

The population must be allowed to discover optimal solutions to presently unknown problems - and thus attain *knowledge* - and I will argue that this is possible only if the system does not punish redundant investigators by depriving them of sustenance, epistemic resources and the possibility of Epistemic Contribution and thus making them subject to “social epistemic exclusion”. Reliable epistemic processes, which I will refer to as “social epistemic inequalities”, are desired - but their reliability is justified only if it withstands contest and disagreement. While the tyrants and the stupid may repress the truth by violence and lies, it will still withstand contest and disagreement every time when challenged in a game of giving and asking for reasons. On the other hand, it cannot be justified if it was not contested. As John Stuart Mill recognized clearly in his epistemic-instrumentalist argument for free speech (Mill 2003), truth must be defended to be knowledge. There is nothing to guarantee that the tyrants and the stupid will not win - and there is no other hedge

against epistemic catastrophe but pluralism and inclusion. Delegating the totality of epistemic labour to a currently convincing group is the strategy least likely to avert the victory of ignorance over learning. The guarantee of universal right to challenge them is the only protection against the enemies of knowledge - they are *least likely* to win if *you* can disagree.

Normative pluralism need not translate into success. It may not in specific cases of reliable social epistemic inequality, in situations when resource is easily defined, the unpredictability low and thus those that already know the solution appear right. When the problem is complex and the epistemic environment dynamic and unpredictable, however, normative pluralism is a better bet from the standpoint of Reason-as-such. *An epistemocrat would do it that way.*

The total search for knowledge is on this account ever incomplete, but particular instances of knowing, the events of knowledge, are however certainly possible - *and pluralism is their condition*. “Locally transient pluralism” will denote the acceptance of a strategy for solving a particular problem which justifies the exclusion from normative community. It is however epistemically instrumental for pluralism to be sustained at the “global” level, spanning the total set of problems, in the form of “universal epistemic inclusion”. This is the fundamental insight in the design of social epistemic systems. The substantial protection of redundant investigators and the freedom to be wrong are thus lessons in positive epistemics.

Institutional Epistemology and Governance: The Infrastructural View

As the climate crisis unfolds in global disorder, inquiries into our capacities for learning and solving complex problems become an urgent concern. Boettke, Tarko and Aligica (2016, 163) stress that the “(t)he question ‘Which collective choice arrangements have the best epistemic properties?’ is one of the most important neglected questions in political economy.” Their research agenda of “comparative institutional analysis based on the collective learning capacities embedded in

alternative institutional arrangements” (Boettke et al 2016, 163) reflects precisely the ambition of IE to find those system-level properties of governance of a large and normatively complex populations which are conducive to knowledge. This thesis shows that IE can recognize redundant normative pluralism and universal inclusion as minimal such properties. Our governance and political economy should, thus, as well. The capacity of a population for production, innovation and problem-solving rests on its capacity to learn. In order to learn, the population must protect redundantly normatively diverse investigators. Markets, liberties and democracy are necessary but insufficient to account for this protection - the access to sustenance and epistemic resources are required for the proper utilization of knowledge dispersed through the population, and more broadly, for the superior governance of collective intelligence.

A secondary purpose of these investigations is then to create a strong core argument in institutional epistemology which can be translated into the urgent and highly relevant new efforts within the broader concerns of governance. In terms of political decision-making and problem-solving with regards to commons, the institutional diversity will outperform a system exhaustively described with a centralized body, as well as with exclusively democracy or markets. Its foundation, which is all I am presently concerned with, must be redundant pluralism and thus universal epistemic inclusion.

As noted in Chapter 1 when discussing the work of Friedrich Hayek, markets, while an immensely valuable epistemic development, endanger this prospect by failing to protect bad betters and thus foreclosing pluralism. As noted in Chapter 1 when discussing the work of John Dewey, democracy, on the other hand, while certainly valuable in epistemic terms as are markets, also harbours epistemic defects - its intelligence harvest mechanisms (voting and deliberation) are flawed, and if not constrained by an additional institutional solutions (for instance, constitutionalism, qualified egalitarianism, polycentrism), it may overextend popular control to violate pluralism and universal inclusion. In this respect, the criticisms of markets and democracy I will provide, on top of being directed at Friedrich Hayek’s and John Dewey’s work specifically, may also extend to positions supporting exclusively market relations (“epistemic liberalism” [Cerovac 2018]) or exclusively democratic

order (“epistemic democrats”) with regards to the epistemic affairs of the population. Epistemological justification for both epistemic liberalism and epistemic democracy is the pluralist project in IE, but they are insufficient for, and may be detrimental to, its proper advancement. I will argue that the superior social epistemic system must, aside from personal and political freedom developed through constitutional liberal democratic order, reflect the pluralist project through developing an institutional guarantee of freedom from poverty. The political and social freedoms are quintessential for any chance of epistemic development or progress. However, without the material and political conditions for Epistemic Contribution, which include guarantee of food, water, energy, shelter, healthcare and epistemic resources, the political liberties are insufficient for the population to be positioned best in its search for knowledge.

The position advanced here can be regarded as Infrastructural View. I will argue that epistemic development and progress is more likely when the population has access to an institutional infrastructure of universal social epistemic inclusion. Liberties, markets and democratic procedures can meet their epistemic potentials only within the appropriate infrastructural settings of universal provision of institutional epistemic inclusion - access to sustenance, epistemic resources and possibility of Epistemic Contribution.

I will argue against strong epistemic perfectionism when discussing Talisse’s proposal concerning the agent development as a comparative standard in IE, in Chapter 2, and retain a scepticism towards any project which would aim to improve the *individual* epistemic virtue *on the level of a population*.

I will not deny normative communities the discretion of promoting their epistemic virtues and enforcing them among their members, nor will I argue against certain epistemic standards within certain epistemic situations. This would be *anti-epistemological*, clearly. All the disagreements should not necessarily last forever, and I claim no such thing. *The condition of freedom to disagree* should be institutionalized through the epistemic infrastructure.

I will here concern myself only with the level of the population, to offer a recounting

of IE's *fundamental* principles in the design and assessment of social epistemic *systems*. There are good solutions in the design of institutions which govern our everyday situations, and I argue that we can discover them only in a system which at least allows us to dissent when they are bad and to pursue alternative strategies. Knowledge must withstand contest. And contest requires *substantial rights of dissent and exit* - not only a formal possibility, but a material and political condition of all individual agents. Thus universal social epistemic inclusion at the level of the system is the true expression of epistemically instrumental redundant pluralism and the proper understanding of the lesson of agent scarcity due to agent suboptimality, which I will discuss in Chapter 2.

It is the institutional arrangement and the minimal shared rules which allow for the attainment of knowledge which is the central target of these investigations, and not individual suboptimality as recognized in the understandings which assess the individual performance according to the idealized idea of rationality (Ostrom 2005, 101; Boettke et al 2013). This is, then, not an argument for an invisible hand approach, but for an attempt at the institutionalization of adversariality through universal epistemic inclusion. This is moreover by no means a relativist account. Knowledge is conditioned on disagreement because it allows a true belief to become justified. Social epistemic inequalities are allowed and justified precisely by resisting system-level exclusion - process must withstand contest to become reliable. It is, lastly, neither an argument for flat epistemology - to say that the redundant investigators must be included does not entail that there should be no social epistemic inequalities, that all Epistemic Contributions are equally valuable or that they should not be disagreed with. Quite the contrary, it is an argument that social epistemic inequalities will be more epistemically virtuous in a system which allows for disagreement through the substantial right of dissent and exit provided through minimal epistemic infrastructure.

Infrastructural View would largely favour investigations into polycentric governance as the development of institutional epistemology and its pluralist project. Redundancy and diversity of institutional solutions, including but irreducible to democratic and market processes, and divorced from the panacea constraints of state-market dichotomy, should be researched and described in IE. Crucially for this

phase of the investigation, there may exist a variety of institutional solutions to the problem of the universal access to sustenance, epistemic resources and the possibility of Epistemic Contribution. IE may recognize certain configurations which give rise to a particularly robust and sustainable epistemic infrastructure - however, for the purposes of retaining openness to discovering ever-more-optimal configurations, the focus of application here will not be on describing the set of best institutional arrangements, but on the tracking of indicators of epistemic infrastructure within a varied, complex and unexpected institutional diversity through the Epistemic Infrastructure Index, as described in Chapter 3. While discussions and research into particular institutional configurations conducive to knowledge is of immense relevance, I am here concerned with the most elementary tool for their comparative analysis. Translated into the principle of governance, Infrastructural View would advance a constitutional guarantee of system-level epistemic inclusion - it should however leave to the population the possibility to experiment with institutional diversity apt at delivering the infrastructure in question. The comparison among the particular institutional configurations, as well as their applicability in a variety of socio-cultural contexts, is an ongoing research task of IE.

My aim is to focus on a particular set of problems - the foundational principles of institutional epistemology. And there *minimal* progress is possible with pluralism and inclusion. Without pluralism and inclusion, the population is more likely to get stuck on a catastrophically stupid idea.

Constitutionalism which along with freedom from oppression guarantees freedom from poverty, and thus contains an *infrastructural amendment*, is required to fix liberalism, democracy and markets in terms of their epistemic properties. Only by protecting the unlikely contributor may we discover Reason.

1. THE PLURALIST PROJECT IN INSTITUTIONAL EPISTEMOLOGY

The present chapter traces the development of the institutional epistemology as founded on the understanding of pluralism as constitutive of superior social epistemic systems.

First part introduces pragmatism as the philosophical school most relevant for the development of basic philosophical commitments in institutional epistemology, and John Dewey's translation of pragmatism's lessons into the experimentalist project in governance of the search for knowledge in a large and complex population of epistemic agents. Democracy as the institutional arrangement most expressive of experimentalism will be discussed, and polycentric governance will be argued to be the more developed option. The role, and limits, of votes and talk as intelligence harvest mechanisms will be examined.

Second part focuses on epistemic suboptimality of epistemic agents and the role of norms in overriding them, thusly outlining the core of Hayekian institutional epistemology. It is then shown that Hayek's foundational insights imply that, over and above decentralization, diversification of normative commitments provides a "hedge" against the tendency of conserving suboptimal strategies inherent to norms. Market as an institutional arrangement and prices as an intelligence harvest mechanism will be discussed, and their limits explicated. Constraining social epistemic exclusion upon bad epistemic "betting" will be recognized as a particularly relevant task of a superior social epistemic system.

Third part is devoted to the contemporary developments in institutional epistemology which feature the use of simulation and agent-based modelling for investigating certain claims in the field. The division of epistemic labour is explored through Weisberg and Muldoon's accounts of mixed groups of agents favouring and avoiding previously successful approaches in the search for knowledge and Zollman's inquiry into conservation of suboptimal epistemic states as epistemically beneficial at the group level. Furthermore, it is shown how Diversity Trumps Ability.

These works are key references in institutional epistemology. While there certainly are authors presently merely noted and not fully represented, the here chosen nodes of the debate are seminal and illuminate the foundational discovery of institutional epistemology - "(i)n any design process where there is substantial probability of error, having redundant teams of designers has repeatedly been shown to have considerable advantage" (Ostrom 2005, 284).

I will now proceed to examine pragmatists', and in particular Dewey's, contribution to the development of institutional epistemology.

1. 1. Pragmatism, Experimentalism and Democracy: Deweyan Institutional Epistemology

1. 1. 1. Pragmatism

Pragmatism lays significant philosophical groundwork for the development of institutional epistemology. Since an excursion into a detailed history of pragmatist thought is beyond the scope and interest of this text, I will focus primarily on the account of theoretical commitments which make pragmatism a relevant philosophical project with regards to the present concerns.

Pragmatism presents an image of the epistemic situation as the one in which an epistemic agent performs an epistemic *action* within an environment which "responds" to the action with a consequence which is then available to the agent to give the initial action a particular epistemic status. The epistemic agent is the one which can offer a judgement (which is a *doing*, an introduction of change into the environment) and respond to its consequences by judging their compatibility with the commitments the agent made by making the initial judgement (Brandom 2001), thus establishing its epistemic status. The epistemic status is awarded within and in accordance with the network of judgements (and thus, beliefs) which the judgement

in question is a part of - in Brandom's vocabulary, the commitment must exhibit material and conceptual compatibility. The epistemic status may be one of truth or falsity, but may as well be one of provisional utility of a belief for further inquiry into the subject matter, of giving more or less weight to a particular belief within a certain network of beliefs, of informed conviction, or, in many cases (as will be investigated in Chapter 2 with regards to epistemic suboptimality of agents), of habitual reassurance even when the consequences of the action do not stack up to form evident and forceful reasons for such reassurance.

This last notion is crucial - it points towards the possibility that epistemic agents may retain beliefs for which there are available evidence of falsity, and thus continue to perform actions which will create consequences available to deem these actions strongly epistemically flawed and which will still not be taken up by the agent as such. Pragmatism thus allows for the commitment (which is always an epistemic *action*) to be either epistemically valuable or void within certain parameters outside of the agent's epistemic "world" (*the space of reasons*), as well as allowing for agents to use the belief in their epistemic "world" *in accordance or in discord* with the value of the claim of the belief outside of that "world" (within the state of affairs). The material incompatibility of a commitment is not sufficient for the epistemic (social) change to take place - a conceptual incompatibility must be accounted for and *committed to by the agents*. Thus pragmatism allows for both attainment of knowledge and for a continuity of ignorance to play the part in the complete account of the epistemic life of a population.

The notion of the epistemic life of a "population" has made its way into this introduction to pragmatism for a specific reason - pragmatism presents the search for knowledge as a *social* activity. As the messy mechanics described above allow for both ruptures of investigative ingenuity and stable choreographies of insensitivity to evidence within an individual epistemic agent, so they do within communities and, as sums of communities, populations. Populations of agents act on beliefs and adjust to consequences of those actions, both theirs and others. Groups (what I will refer to as normative communities) act on same sets of beliefs, and sometimes agents within those groups form beliefs different from the group. Sometimes this is because certain agents within those groups respond to evidence the other members

of the groups have become habitually unresponsive to (Gronow 2012, 28-31) This last claim points to the need for clarification of another aspect of the sophisticated complexity of the pragmatist account of the search for knowledge - its *inferential* nature.

The social aspect of the search for knowledge is crucial because it allows us to paint a picture of numerous agents searching for knowledge by use of group knowledge (very roughly put, what the groups have taught it to believe) but also by the use of a unique inferential apparatus each epistemic agent possesses. The inferential nature of knowledge is key to understanding it as an action. Without the concept of “inference” pragmatism cannot resolve the gap between a belief and evidence which would allow for the belief to be updated by evidence or *not* be updated by the same evidence. Agent is active *and epistemic* precisely in this moment - when the belief as something with which the agent “comes into” a situation and evidence as something that “is there” in the situation “clash” and *feed back* into a new epistemic state within the agent. Either a new belief is formed or the old one retained. This depends on a particular inferential activity of the agent, which is by no means wholly under the agent’s control but is nevertheless (in a strong sense) unpredictably unique to that agent.

The populations of agents thus form beliefs and encounter evidence, and each agent adapts its beliefs to the encountered evidence in a particular manner, thus changing the future trajectories of the search for knowledge of the populations. Within this image, pragmatism introduces another concept crucial for understanding why the inferential uniqueness does not lead to outright epistemically distortive chaos among the epistemic agents - the concept of *habit*, which I will later in the text subsume under the concept of norm.

Inferential practices are to a relevant degree influenced by habits of the community an agent is a member of. Does an agent update a belief in the encounter with the new evidence or does it retain the old belief despite the new evidence is not only a result of the evolutionary roulette of pre-social inclinations and biological fluctuations which may be introduced as character traits or other cognitive or psychological variations within the population - but is also a result of socially-incentivized available

patterns of recognition and weight-distribution with regards to evidence as state of affairs which needs to become a variable within the inferential machinery used in the search for knowledge. These patterns are *norms* of recognizing a certain state of affairs *as evidence*, of giving it salience and thus conceptual “shape” at least to a degree that it can be used in the epistemic activity (translating it into a judgement, and further, *a reason*), and of giving this state of affairs recognized as evidence a particular weight (or, a particular relevance) within the network of commitments available to the agent. These norms guide the inferential practice - and while unpredictably unique inferential practices (and thus an unpredictably unique set of norms) are a relevant aspect of each agent and of epistemic life of populations, a *significant* amount of norms that guide an inferential practice are distributed and enforced by the community the agent is a member of.

Judgement, belief and norm are considered in pragmatism as somewhat akin to primary units of epistemic performance primarily because they are conceptual contents which *can* be made explicit in order for their success to be evaluated and thus can constitute *epistemic* action, one which exhibits capability of undergoing revision upon proper evidence. This *does not* mean that there are no implicit judgements, beliefs and norms - *the most of them may as well be*. What it does mean is that their common trait is that if they cannot designate *a condition of the success* of the judgement, if they cannot be tested and contested and examined for their material and conceptual compatibility, and thus if they cannot be *assertible*, they cannot be called epistemic at all - their core function cannot be search for knowledge. They may play into this search in some indirect way (the way all non-epistemic conditions, for instance availability of food and shelter, may play the role in epistemic performance), but until they have become something with which one can *do* something in the environment and by doing *commit* to a particular epistemic evaluation of that doing and thus bring about change in the environment which *can be* epistemically evaluated, they are not epistemic. They need not in full actuality be contested, nor do their conditions of success need be actually responded to properly by any agent - it is the *availability* of this, as Brandom would have it, *deontic scorekeeping*, which counts.

“On the side of the *consequences* of acquisition of practical deontic statuses, it appears in the essential role that propositional, that is, assertible, contents play in specifying conditions of *success*: that is, what would count as fulfilling a commitment to act. Forming an intention (undertaking a commitment) to put a ball through a hoop requires knowing what it is to put a ball through a hoop—what must be *true* for that intention to *succeed*.”

(Brandom 1998, 4)

It is important not to see pragmatism as relativism of the sorts which forgoes the possibility of attainment of true belief. It is the crucial aspect of pragmatism that true belief is possible, because otherwise inference wouldn't be (given that the notion of inference relies on updating of a belief according to a certain state of affairs, and thus depends on the reference to a proper and improper reaction to a certain state of affairs), a distinction between a better or worse habit (and norm) would be seriously shaken (because there could be no reference to better or worse patterns of inferential practice without some form of knowledge discovery or production), and the enterprise of search for knowledge as a social activity wouldn't be possible (because there would be no knowledge to be attained, and no social organization could claim to be in a better shape to attain it). Moreover, at the meta-level of the discussion, the very image of the search for knowledge endorsed by pragmatism is “complicated” in the manner described above precisely because *it is closer to a true image* of search for knowledge.

This is how agents know or remain (or become in a different way) ignorant - by making use of an incomplete, partially coherent and, as will later be discussed, only somewhat controllable network of rules for updating or retaining beliefs in the active encounter with the world outside. This makes them *epistemic* agents. The cognitive content and the inferential practices (in some cases socially-enforced and in some cases individual-specific) are active “machines” for navigating the landscape of the unknown and the unpredictable (and, usually, problematic) states of affairs, and the signals from the landscape feed back into the “machines” with variable relevance and force and in particular conceptual “shape” of variable usefulness for further navigation.

Given this, it would be useful to highlight that the concept of particular relevance within pragmatism could be said to be the ability to *recognize error in judgement* (Brandom 2001). While epistemic agents may retain a false belief despite evidence to the contrary being available both as a mere state of affairs the agent has not even, so to speak, “taken into account” or *conceptualized*, as well as a conceptualized “item” available for further navigation of the epistemic landscape but neglected, and while, as I stressed a bit, the agents may even *tend* to do this more frequently than not, it is important to stress that they may also change their mind, revise an erroneous commitment.

This is a peculiar trait , and one that I will claim in Chapter 2 to be a prime candidate for the comparative standard in the design and assessment of social epistemic systems, and thus in the case of population-level account of search for knowledge. As it is with populations, so it is with agents - their epistemic quality depends on how likely they are to revise erroneous beliefs¹.

The proposition of revision of erroneous belief, or for that manner, normative commitment as a central epistemic good points to what could be called pragmatism’s central epistemological claim that social epistemic progress and development consists of *moving away from untruth*. As I will argue in the Chapter 2, shifting away from the attainment of truth as the comparative standard towards the ability of the social epistemic system to revise the suboptimal epistemic state is crucial for establishing how likely is a population to actually find the truth.

The history of pragmatism is a long and complex one, and there are certain disagreements between pragmatists I will overlook here. Moreover, certain philosophers may be recognized as more relevant in their contribution with regards to the finer accounts of the inferential and normative nature of epistemic agency and knowledge (the present text will largely derive the deeper philosophical claims in this

¹ As things will complicate further, there will in Chapter 2 be significant talk of how some populations epistemically perform better when (at least some) agents perform conservatively - by preserving erroneous beliefs to a certain degree. From the perspective of each agent it is better if it were to change its mind more efficiently. This, I will claim, is where the key for disregarding something I will call agent development (and includes virtues of individuals) as a comparative standard of design and assessment of social epistemic systems lays - the better social epistemic system cannot rely on the development of individual epistemic agents because what is best for the agent may not be best for the system. What is best for the system, however is best for the agent because it elevates the overall epistemic quality of all epistemic materials and inferential practices (normative commitments) available to the epistemic agent.

regard from Sellars and particularly Brandom). But Dewey notably understood the relevance of the revision of suboptimal epistemic state for the design of social epistemic system, pursuing *experimentalism* as the principle for choosing an institutional arrangement governing the population of epistemic agents.

1. 1. 2. Experimentalism

The insistence on revision (and inference) which pragmatism builds on holds a particular instruction for epistemic agents to test and contest their beliefs and commitments, even in the circumstances of never being fully certain of any of their tests and contests being valid. Experimentalism is the crucial part of pragmatist epistemology. In line with John Stuart Mill's understanding of knowledge (Mill 2003), pragmatism purports that prior to the action of testing and contesting a belief, the belief cannot be regarded as knowledge at all. Furthermore, in numerous cases, our epistemic habits will make us incapable of recognizing the error in judgement. Only within the population of epistemic agents which freely test and contest commitments, and thus under the conditions of *pressure of free normative pluralism* (in which disagreement is not stifled but utilized as intelligence), do we gain a slim chance at revising our suboptimal epistemic state. Dewey in particular recognizes this, and makes an *institutional* claim of experimentalism - the political epistemic good, the knowledge on how to govern best, is conditioned on the policies being tested, contested and revised upon detected failure (Dewey 1976; Dewey 1977, 15; Anderson 2006, 13). To allow the contest of the commitment on how society is to be organized *is* the "best bet" for making the society best organized.

The descriptive and prescriptive claim of normative pluralism, must be distinguished here. Pluralism may be understood as a given (Gaus 2018a). The populations of any number of agents contain a certain normative pluralism derived from a more fundamental *inferential* pluralism (or what Landemore [2012] refers to as "cognitive diversity") I earlier pointed out as a large pragmatist concern. Two agents are already normatively diverse if there are two, and the more there are, the more complex the

diversity becomes - it becomes a diversity of normative *communities*, meaning basically groups of a certain number of agents destined to live in the world of agents with which they *do not* share normative commitments. There is ample evidence that human epistemic agents live in substantial epistemic disagreement, both among communities and in communities. This is, both Dewey and Rawls would agree, how things are. The large populations will always be pluralist - normatively complex. The particular prescriptive claim would extend that pluralism is epistemically beneficial. The pressure to change a suboptimal habit can be derived *only* from agents and communities *freely* testing and contesting a *variety* of normative commitments - only the freedom to test and contest, disagree and dissent, allows for the error to be recognized and for the population to *learn*. Moreover, given the environment is dynamic, unpredictable and abounds with contingencies (and it must be clearly understood that *the environment includes other epistemic agents*), the need for continued testing of beliefs and commitments (epistemic actions) amplifies - the circumstances change, and the habits better change accordingly.

Normative pluralism (even redundant) “piggybacks” on experimentalism - reducing found pluralism reduces experimentalist developments, given it effectively shuts down the contest of commitment, and thus disagreement. Properly understood the better norm can become available *exclusively* if norms can be tested at all. Talisse’s work is of particular precision when he distills this understanding to its essential statement - without the ability to contest a commitment, an agent cannot even be regarded as epistemic at all (Talisse 2009). To experiment with and contest commitments is to be an epistemic agent. Furthermore, from the standpoint of benefit for the large population as many normative communities as plausible should be allowed and encouraged to test and contest their normative epistemic states and moreover the exposure, interaction, cooperation and competition among them should be encouraged - this is the only way the failure of an erroneous normative commitments can become *evident*.

Thus, the recognition of the error in judgement a central epistemic good for pragmatism, and the superior political epistemic system is *experimentalist*. To recognize and revise failed normative commitments on every level, and thus also on the level of *policies and institutional arrangements*, is to *learn*, and thus progress in

the search for knowledge. Dewey believed the principle of experimentalism is best expressed by *democracy*.

1. 1. 3. Democracy

Democracy gives the epistemic agents a chance to experiment with policies at different levels of governance and most importantly *may foster relatively* (among other competing systems) *most efficient adaptations* at the recognition of policies' failures (Dewey 1977, 19). I will give an account of Dewey's democracy, particularly relying on Anderson's understanding of it (Anderson 2006), and proceed to argue that the further inquiry into experimentalism as a governance principle can be recognized in a system of multi-level institutional diversity featured in Ostrom's work on polycentric governance.

The concerns about the epistemic quality of democracy have recently become a relevant and animated field of research. The accounts and positions, frequently referenced with regards to particular arguments in this text, abound (Estlund 2008, Talisse 2009, Landemore 2013). Elizabeth Anderson, in her seminal paper "The Epistemology of Democracy" (2006), puts forth an argument that, when pitted against Jury Condorcet Theorem and Diversity Trumps Ability Theorem, Dewey's experimentalist model of democracy is most apt at utilizing the dispersed knowledge in the population.

"Most importantly, Dewey's experimentalist model of democracy helps us see the epistemic import of several democratic institutions that sustain its dynamism, its capacity for change: periodic elections, a free press skeptical of state power, petitions to government, public opinion polling, protests, public comment on proposed regulations of administrative agencies. In Dewey's model, these are mechanisms of feedback and accountability that function to institutionalize fallibilism and an experimental attitude with respect to state policies."

(Anderson 2006, 14)

The central finding of institutional epistemology is Hayek's claim that the information required to solve a particular social problem is distributed unpredictably in the population (that *knowledge is dispersed*). Furthermore, Hayek claims, and Anderson requires of IE to face, that no central body can adequately gather nor use the information, and knowledge, dispersed in the population. What follows is *the need for the system which can harvest the knowledge and streamline the problem-solving most efficiently*. For Anderson, Dewey's understanding of the epistemic functions of democratic procedures for harvesting the information, and especially the possibility it leaves for dissent to take place even after the collective decision has been made (a solution to the problem accepted), is what is needed for the system to be experimentalist, to allow for the testing of a policy, for the consequences of a decision to be recognized as evidence to its success or failure, and the decision given epistemic status in this light.

“Without an opposition to remind the public of continuing objections to collective decisions, and to pose alternatives, accountability of decision makers is impossible. Nothing would force decision makers to reconsider their decisions. Only with such continuing opposition can fallibilism and the institutional capacity for experimentation—revising one's decisions on the basis of experience with their consequences—be realized.”

(Anderson 2006, 17)

Condorcet's Jury Theorem, as I will expand on when discussing votes and talk as intelligence harvest mechanism, has a strong problem with accounting for votes as primary vehicle of democratic epistemic value because in order to do this it must imagine the agents as probably right in making the decision - while institutional epistemology should regulatively count agents as epistemically suboptimal, and thus probably wrong. Dewey's experimentalism requires them to be *able* to learn *whether they were right* and make a better decision (allowing they may not). Diversity Trumps Ability (DTA), which I will also expand on later, while capturing the value and the relevance of disagreement (and redundant normative pluralism) within a population, does not appropriately account for particular democratic mechanisms for pooling,

gathering and streamlining information from the population, what I will later refer to as *intelligence harvest mechanisms*. This is predominantly because DTA is concerned with other matters - but as an immensely valuable contribution to institutional epistemology itself, it would surely warmly welcome findings on the intelligence harvest systems into the field. Anderson particularly stresses that Dewey's experimentalist account did, however, model the harvest just right - in the trinity of harvest in the institutional epistemology, prices-votes-talk, Dewey's expansion of possibility for the universal inclusion into the problem-solving of governance, and particularly the possibility to express dissent even after the decision has been made, accounts for, roughly put, votes with *feedback* and talk with *reasons* (Anderson 2006, 14).

Dewey understood, however, that both of these may be hijacked by a problematic, malign, group, and thus instructed the citizens adopt an *ethos* of democracy in order to successfully sustain its experimentalist epistemic value. The problem with the internalization of an ethos by a population, even if it were not regarded as a practically unsound policy (particularly given the found pluralism of large populations), is the deeper understanding that an adoption of a certain seemingly epistemically favourable trait by the agents need not result in overall epistemically best performance by the population (something I will expand on in the section of agent development as comparative standard in institutional epistemology).

Democracy, moreover, is an order open to epistemic distortion primarily inasmuch as the intelligence harvest mechanism by which it operates may fail or be overextended to particular areas of potential problem space in which *popular control impedes Epistemic Contribution of the individual epistemic agent*. This second problem may likely appear when democracy becomes a highly centralized governance model in which the average becomes the center. Popular will itself, if used improperly, may become an institution which highly overestimates its epistemic scope or simply mistakes its epistemic domain. Moreover, media as the telecom of reasons, the space of talk, may become epistemically corrupt inasmuch as it may frame the common task or insist on a dissent erroneously, at times possibly in favour of a normative community highly malignant to the epistemic cooperation. Democracy may, then, foreclose itself to pluralism (Kelly 2006).

Additional institutional framework appears required for democracy to sustain itself without derailing into epistemic distortion. A set of constitutional liberties appears to be the most effective primary safeguard to such tendencies. Polycentric governance, in which self-management, higher level governance and multiple providers of institutional services (a diversification of institutions) “work together” to rearrange the epistemically favourable democratic traits into variously less suboptimal configurations, appears to contribute to staving off of democracies’ more dire excesses.

All of this said, however, democracy is most certainly among the most valuable epistemic discoveries of human civilization. Its various multi-institutional instances have severely restrained the most anti-epistemic institutional practices of centralization, and the voice of the people *can* silence the oppressors. All the Dewey’s favourite democratic traits even at their worst, and I believe even when they may *appear* extremely problematic to the epistemic agents themselves (as they troublesomely do now [Foa and Mounk 2016]), still surely perform better against numerous alternative proposals of centralization. They may, however, be upgraded. And it is this upgrade that is the proper issue of the design and assessment of social epistemic systems - the particular institutional ecology within which democracy should fit to contribute to the development of the social epistemic system which is more likely to produce the highest epistemic good, lead the population *farthest away* from worst kinds of ignorance.

1. 1. 3. 1. Polycentric Governance

By both reducing the scope of plausible management of ethos-internalization to particular groups within a larger system of groups, and, more crucially, by diversifying the institutional ecology, polycentric governance could be regarded as a more refined recent development of experimentalism in political economy and institutional epistemology. As I will show in the next section, votes and talk are tools

of immense value, but by themselves and without specific qualifications insufficient to perform the most “fruitful” harvest or guarantee a sufficient hedge against consistently suboptimal decision. Moreover, Dewey’s particular democratic vehicles for votes and talks work along the lines of citizens communicating with the state in order to influence the policy-making. While state most certainly should remain open to communication and engagement with the citizens in decision-making, it appears a more decentralized and diversified institutional ecology may be required for proper account and design of the intelligence harvest mechanisms within the social epistemic system.

First concerned with governing the common resources and *self-organized* collective action (Ostrom 1990, 7), Ostrom continued to worked on the development of “a grammar of institutional diversity” in order to understand how is it that agents of severely limited epistemic capabilities encountering an environment of dynamic production of severe problems *may* perform epistemically optimally, and which institutional setting, or more precisely which arrangement of rules involved in an action situation at which level of problem-solving, is most conducive to this result (Ostrom 2005; Ostrom 2010, 436) The framework developed by Ostrom and her co-researchers in polycentric governance allows for a finer effort at the articulation of the systems of nested and multi-level problem-solving normative communities, as well as for particular investigations into rules governing groups with more chance of epistemic success. For present purposes, however, polycentrism makes a relevant move in the experimentalist account of a social epistemic system by *diversifying the institutions*, both descriptively and normatively.

Key to understanding how Ostrom’s highly decentralized system fits the experimentalist image is in its treatment in error (policy failure) and conflict (disagreement). By diversifying institutions, polycentrism 1) allows for conflicts, the disagreement between units, to become conducive to epistemic benefit because they stimulate the communication of *unique information* which would most likely be held implicitly by the units lacking the conditions of disagreement (Ostrom 2005,

286)², and 2) allows for errors, the evidently failed policies, to be most *localized* and thus least threatening to the stability of the large-scale system, and thus *more likely to become evident* then in the case of system-level failure of policy, where the immensity of consequences will likely lead to wholesale new bundle of more urgent and wicked problems (Ostrom 2005, 284)³ which tend to drive communities to stronger conservation (Heinrich 2009). These accommodations of error and disagreement within the account show how closely Ostrom's work can be seen as related to the pragmatist, and particularly the experimentalist, project, and how valuable both are in the investigations on the design and assessment of social epistemic systems.

Dewey saw the experimentalism embodied in the democratic institutions of voting and free disagreement and in agents adopting an ethos of experimentation, Ostrom offered a complication by developing a sophisticated analysis of how complex multi-level institutional diversity of largely autonomous and self-organized units "may be more effective in learning from experimentation than a single central authority." (Ostrom 2005, 218). Given numerous overlaps of theoretical interests, and in particular new institutionalism's (Ostrom 1990, 29) theoretical debt to Dewey's and pragmatists' work, I believe observing their efforts as a continuity, particularly within the above-sketched outlines concerned with the design and assessment of social epistemic systems, is quite inviting (Aligica 2014, 66, chapter 6).

² "Polycentric systems can generate considerable conflict among the various units at multiple levels due to their interdependence. Conflicts that escalate from misunderstandings to ever more serious charges and countercharges that turn to violence are certainly negative processes. Conflict may, on the other hand, generate more information that is useful to participants in their efforts to solve challenging problems. Ebbin (2002, 2004) has traced the evolution of conflict in the comanaged salmon fisheries along the coast of Washington both as a fishery biologist working with several of the tribal organizations and as a researcher conducting extensive interviews with participants at all levels. While the early conflict was framed as a technical problem regarding the knowledge to be used in managing the system, later conflict "focused on questions of equity and conservation as well as authority and jurisdiction" (Ebbin 2004, 82). The redefinition of conflict in the court system and in other arenas helped to create new institutional mechanisms that "changed the rules of the game and the processes in which new conflicts are addressed" (82). After some experience with the new institutions, even the government officials recognized that new information was being generated that initially led to more conflict but eventually led to better management of the stock." (Ostrom 2005, 286)

³ "When there is only a single governing authority, policymakers have to experiment simultaneously with *all* of the common-pool resources within their jurisdiction with each policy change. And, once a major change has been made and implemented, further changes will not be made rapidly. The process of experimentation will usually be slow, and information about results may be contradictory and difficult to interpret. Thus, an experiment that is based on erroneous data about one key structural variable or one false assumption about how actors will react can lead to a very large disaster (see Wilson, Low, et al. 1999). In any design process where there is substantial probability of error, having redundant teams of designers has repeatedly been shown to have considerable advantage (see Landau 1969, 1973; Bendor 1985). The important point is: If the systems are relatively separable, allocating responsibility for experimenting with rules will not avoid failure, but will drastically reduce the probability of immense failures for an entire region." (Ostrom 2005, 284)

What is most relevant to the present inquiry is, however, Dewey's central finding that the institutional arrangement capable of producing the greatest epistemic value must be experimentalist. The social epistemic system must *learn*. The particular institutional order most apt to learn, however, only begins to be built by democratic institutions - its optimal shape, if there is such a thing, is yet to be fully described.

I will now offer a further comment on votes and talk as transnormative intelligence harvesting mechanisms.

1. 1. 4. Votes and Talk

The dispersed knowledge in a large and normatively complex population must be communicated across and among diverse normative communities. The means of such a *transnormative* signalling of epistemic content conducive to the harvesting of collective knowledge are "intelligence harvest mechanism", of which institutional epistemology recognizes three broad traditional categories - votes, talk and prices (Anderson 2006). This section will address votes and talk, intelligence harvest mechanisms representative of the epistemic properties of democracy, while prices will be accounted for in the course of discussing Hayek's contributions to institutional epistemology.

Democracy is the institutional order apt at delivering the experimentalist program, and the primary means of transnormative communication in it are votes and talk. Regarding votes, their basic epistemic value lies in the task of harvesting the knowledge dispersed in the population (as will be explicated, through *differently* erroneous commitments "cancelling each other out") by giving governing powers to citizens in order to defend the system from being hijacked by a certain epistemic elite. Votes are informative through transforming the signals from the population into an output through averaging the Epistemic Contributions, and in democracy their binding powers have an epistemic function (which they may fail to perform, both due to broader epistemic suboptimality of agents and to the more specific epistemic

distortions of democracy). Regarding talk, after arranging themselves in groups with common rules or interests, citizens use media to provide each other with information and reasons. Crucially for Dewey, even after votes “seal the deal” through averaging, talk allows the citizens to object to the output, thus allowing for the new adaptive trajectories to emerge, with feedback on the output potentially reframing the issue and allowing for the change in future voting preferences to take place through the repeated encounter with dissenting views and their reasons.

This image of votes and talk is fundamental in the sense that it justifies their elementary utility with robustness and clarity, and in the sense that it obviously lacks the complexity of the real-world situation of votes and talk. This second characteristic can be taken as its flaw, but only if one were either to discuss them in order to garner a more sophisticated understanding of the contingent constraints they feature as particular intelligence harvest mechanisms or to discuss them as candidates for the exclusive intelligence harvest mechanism. Given the latter, I believe, cannot be reasonable proposition for *any* intelligence harvest mechanism, the following discussion is meant to illuminate the particularities of votes and talk without pushing for them as an exclusive panacea for the problem of transnormative communication or intelligence harvesting. They are one of many - and none can be the only.

Condorcet Jury Theorem, a frequently addressed account of collective intelligence, requires that all agents involved in the voting are at least more than 50% more likely to be correct for the aggregation of their votes to be *increasingly* more likely to be correct (Condorcet 1976). This assumption of the smallest likelihood to be correct renders Jury Theorem of limited utility for our present concerns - the baseline condition in play for the design and assessment of social epistemic systems is that agents are suboptimal, and thus more likely to be wrong. When this is the case, the aggregation of their contributions, votes, will lead to increase in likelihood of incorrectness (Sunstein 2006, 33). Thus, following Jury Theorem, suboptimal agents when engaged in en masse intelligence endeavour can only deliver a catastrophically suboptimal output. The utility of thusly understood Jury Theorem for institutional epistemology is diminished - even though the discipline is certainly interested in understanding the features of the system which can accommodate an unlikely optimal epistemic bet from a variety of epistemic agents, its primary concern

is how to mitigate and, eventually, utilize the agents' epistemic suboptimalities. Another attempt at following Jury Theorem in discussing collective intelligence would be to pick only those agents which can be said to be probably more likely to be correct than not - thus leading us into territory of governance by experts. As I will argue in more detail, the problem with such a policy is that it systematically forecloses the possibility of contesting a suboptimal epistemic state because it delegates the epistemic labour to a single normative community - those deemed (by the population or other experts) more likely to attain true belief. When they get stuck, there is nothing to make them recognize the suboptimal state or shake them out of it. It is crucial to understand that these agents, experts, are not optimal - they are, in the best case, less suboptimal. The mystery, then, is how is epistemic progress possible *even* in the population of exclusively suboptimal epistemic agents. As I will show, only pluralism, and subsequently inclusion, can *begin to* mitigate these unfavourable circumstances.

This trouble with Jury Theorem, however, does not exclude votes from the pool of legitimate and, more to the point, valuable intelligence harvest mechanisms. Voting in democratic societies serves a particular function by pooling exactly "random and symmetrically distributed errors" (Schwartzberg 2015, 196) in a Miracle of Aggregation. Sunstein's account of predictive markets presents them as a particularly developed intelligence harvesting mechanism, and voting in this case may allow for significant levels of epistemic excellence if provided proper conditions, among which the most relevant appears to be the investment in being correct. Sunstein claims *the investment in correctness* appears to be among the key features of voting as conducive to epistemic good because by having an explicit stake in the problem being solved, the agents make the most thorough and responsible choice of the bet they believe to be the best, thus effectively making the best bet they can, which is more likely to be more than 50% right (Sunstein 2006, 121). This finding echoes Ostrom's that one of the design principles of successful common resource governance is that those who are affected by the regime should be authorized to participate in modifying its rules (Ostrom 2005, 263-265). The investment in finding the solution by all agents involved in looking for it is a significant feature of groups which tend to eventually find the optimal one.

In numerous cases, votes appear to demand a normative overlap, or a normative “pidgin” (Muldoon 2013, 124), with regards to the understanding or “framing” of the problem, which can be difficult to find in the normatively pluralist society. But, more problematically, the effort at finding the normative overlap required for epistemically most sound voting structures may lead to settling for a particular community’s normative strategy at the expense of others’, which in turn sows distrust in the pluralist social system and, more to the point, may be an epistemically distortive form of hijacking power in the social system. Furthermore, the particular tension between the need for common ground for votes to work optimally and their role as a tool of *transnormative* communication may emerge as problematic within specific situations - fundamentally, for votes to work it requires the voters not to err in the *same* way, and the eventual framing is vulnerable to containing implicit norms (unknown to any agent involved) which may hinder the required pluralism. This need not be the case, but at certain voting situations may. The present discussion points in particular to the limits of the votes’ *robustness* - they must be used with special care as to their appropriateness for the problem situation.

Deliberation is an intelligence harvesting mechanism of similarly complicated nature with regards to epistemic instrumentalism. It may be of particular value for the very task of defining a clear common understanding of the problem and subsequently collective recognition of the solution, in cases when these are epistemically virtuous features of the situation. As Ostrom’s works shows, managing common resources can be performed significantly better if “cheap talk” precedes it, in order for the settling of the “common ground” to take place and in order for understanding one’s interlocutor as invested in finding the right solution (Ostrom 2005).

Sunstein, on the other hand, shows how deliberating groups, when lacking *strong and persistent diversity*, can easily slip into a number epistemically distortive practices - they can settle for the common ground which does not reflect the relevant aspects of the situation out of the need to rein in their dissent or unique Epistemic Contribution for purposes of, roughly put, remaining socially accepted in the group. Or they can, if diversity is particularly lacking, slide right to the epistemically more distortive extreme version of their beliefs, again in response to (even implicit) social pressures (Sunstein 2006). Sunstein’s view of successful deliberation overlaps with

the majority of authors' whose contemporary understanding of pluralism is instrumental - the best deliberation requires a sufficiently diverse group of deliberators (Sunstein 2006, 29)

Practices of public deliberation are a relevant topic in social epistemology as well as contemporary political philosophy. The epistemological understanding of deliberation's standing and relevance within justification of democracy would certainly purport that democracy being the institutional order which can accommodate deliberation most certainly *does* give it an epistemic edge over institutional orders which do *not*. This is simply because no order which excludes the free exchange of reasons among epistemic agents can make a sound bet as to epistemic superiority.

The understanding why this is so can be traced to Mill's account of knowledge as possible only if exchange of reasons can take place (Mill 2003), and it has recently been most soundly defended by Talisse's argument that there cannot be epistemic agents prior to the act of free exchange of reasons (Talisse 2009). Talisse's argument does not rest on the particular epistemic value of deliberation as real-world practice, but it presents the exchange of reasons as *constitutive* of epistemic agents. Thus, instead of putting public deliberation against some or other form of epistemic elitism, it sets the deliberative *action* as the ground of having the population of epistemic agents *at all*, and thus simply removes all institutional arrangements which do not allow for the free exchange of reasons from the competition altogether. They simply do not treat the population as an epistemic resource, and in turn the population does not engage in the epistemic activities at all (in this highly, of course, reduced version of real-world human populations, which tend to engage in reason-exchange in numerous ways as well as when under clear threat if they do). Talisse's account, following the pragmatist program closely, makes it clear that deliberation is a given in the population of epistemic agents. Democracy is "merely" the institutional order responsive to this state of things.

So, as it is to knowledge, deliberation is constitutive to epistemic agency - and thus even if certain deliberative occasions and groups fail, this cannot be of use as the argument against deliberation. These failures are insights into particularities of the

optimal design of deliberative environment and conditions for each particular problem situation. The instructions for customizing according to the task abound in literature, and should by all means be thoroughly consulted and developed in the work in the discipline.

Both votes and talk are, obviously, tools for harvesting collective intelligence in a social epistemic system of immense value. When understood more fundamentally, both are intelligence harvest mechanisms which no social epistemic system could exist without. When finely institutionally structured, they are impressive machines for gathering and streamlining the knowledge dispersed in the population. And the institutional order defined by them is democracy. Just as they are, so democracy appears, *necessary but insufficient* condition for the superior production of epistemic good, the proper epistemic governance.

1. 1. 5. Conclusion

Dewey's pragmatism instructed experimentalism which he used to justify democracy. The understanding that the recognition of error in judgement must instruct the design of the institutional arrangement governing the search for knowledge of a large and complex population is among the fundamental insights of IE. This recognition is conditioned on the contest of the commitment. An institutional arrangement favourable to dissent, and *redundant normative pluralism*, is more likely to attain knowledge, by moving away from the untruth through contest.

I will now move on to the elementary contributions to IE by Friedrich August Hayek.

1. 2. Ignorance, Norms and Instrumental Pluralism: Hayekian Institutional Epistemology⁴

1. 2. 1. Introduction

Among the key contributions of Friedrich A. Hayek to institutional epistemology are his two fundamental claims - 1) that ignorance is a constitutional feature of individuals and groups involved in the search for knowledge, and 2) that individuals and groups search for knowledge guided by norms. My aim in what follows is twofold:

1) to show that these two fundamental claims entail that the population in which individuals and groups follow different, including redundant, norms is less likely to get “stuck” following an inferior path in the search for knowledge - more specifically, that *normative pluralism is epistemically instrumental*;

2) to show that markets are epistemically defective institutional arrangements because they may allow a reduction of the number of individuals following different, including redundant, norms in the search for knowledge - more specifically, that markets violate the principle of epistemically instrumental normative pluralism by failing to guarantee *universal epistemic inclusion*.

I will thus offer a reconstruction of Hayek’s thought in institutional epistemology and a reassessment of Hayek’s favoured social epistemic system, the market, in terms of its epistemic properties. It will show that Hayek’s fundamental claims in institutional epistemology lead to particular findings absent from his account - namely that if we are all ignorant, we should protect the more ignorant among us in the search for knowledge.

Sections 2 and 3 provide a detailed exposition of Hayek’s fundamental claims - section 2 focuses on the claim that epistemic agents (exemplified here by human

⁴ A shorter version of this subchapter has been published as Zubčić, ML. forthcoming. Ignorance, norms and instrumental pluralism: Hayekian institutional epistemology. *Synthese*, DOI: 10.1007/s11229-019-02420-5

individuals) are constitutionally ignorant, and section 3 on the claim that epistemic agents follow norms in their search for knowledge. Section 4 will argue that populations consisting of epistemic agents who follow different, including redundant, norms are more likely to avoid getting “stuck” at following inferior norms in their search for knowledge. Sections 2, 3 and 4 thus contend the following. Hayek understood that “ignorance” is constitutional of epistemic agents, and that only through the decentralization of plans is the search for knowledge made possible. However, the decentralization of epistemic activity has epistemic value only if it entails the *diversification* of normative strategies. If a single plan or strategy were followed by all agents, they would have no protection against being “stuck” on an inferior path in their search for knowledge. It is only if the agents are *differently* suboptimal, if they use different suboptimal normative strategies, *if they do not all keep making the same error*, that the population *may not get* “stuck” at the same poor judgement, false belief, or bad decision. I will therefore claim that Hayek’s primary lesson of constitutional ignorance of normative agents implies that normative pluralism is conducive to epistemic good. This understanding is compatible with foundational accounts in institutional epistemology which show that groups of diverse investigators outperform groups of the most able or individually epistemically virtuous investigators (Mayo-Wilson et al. 2011; Page 2008). It is also in line with pragmatists’ and Mill’s understanding that knowledge is possible only under conditions of contest within the “space of reasons” (Mill 2003; Brandom 2001).

Section 5 will provide an account of prices as means of communication among normatively diverse communities and thus as an intelligence harvest mechanism. Section 6 offers an argument that the market, Hayek’s preferred institutional arrangement of epistemic governance over a large and complex population, violates normative pluralism. It will be shown that this is so because the market lacks what I will define as “the institutional protection of (individual) epistemic agents from social epistemic exclusion upon bad epistemic betting”, which is required for normative pluralism to be satisfied at the level of a social epistemic system. Since the epistemic situation is the one in which “no action is guaranteed to succeed or fail, and no history determines an optimal action with certainty” (Mayo-Wilson et al. 2011, 662), epistemic actions under constitutional ignorance are to be considered “epistemic betting”. In populations governed exclusively by markets, the epistemic agents

making bad bets (and thus following relatively suboptimal and redundant norms) *may* become indefinitely deprived of sustenance, free access to epistemic resources and possibility of Epistemic Contribution (Fricker 2015). The protection against this predicament will be argued to be of critical importance for the design of the superior social epistemic system.

Lastly, I will now reiterate certain terminological clarifications, and furthermore make explicit certain conceptual commitments.

“Epistemic agents” will here be exemplified by human individuals. It will be argued that individuals follow “norms” in their search for knowledge (discussed in section 3), and those groups of individuals that follow the same norms will be referred to as “normative communities”. The baseline epistemic agency is recognized here as a property of individuals primarily because they can infer consequences of norm-following and provide feedback on the success of a norm to the community and the population. Following Brandom, epistemic agents must be able to make a *judgement* and a normative *commitment* which they can *make explicit* and for which they can be held *responsible* (Brandom 2011). Individuals are the source of unique information and interpretation (Sunstein 2006; Fricker 2015), provide epistemically valuable cognitive diversity within groups (Landemore 2012a) and take the risk of changing a norm required for it to change at the level of communities and populations (Bicchieri 2017). I will not present detailed arguments regarding these points here nor discuss group agents.

“Epistemic suboptimality” will frequently be used interchangeably with the term “ignorance” as it appears to be a more appropriate technical term to *include* Hayek’s specific notion of “ignorance” (discussed in section 2). Epistemic agents will be regarded as epistemically suboptimal if the following criteria is satisfied - a) they don’t have access to all relevant evidence, b) they have limited and lacking conceptual resources, c) they make inferential mistakes and errors (systematic mistakes), d) they tend to conserve suboptimal strategies in the search for knowledge, and e) they cannot predict the future. Given that the “epistemic suboptimality” will be recognized as the necessary feature of all epistemic agents,

their epistemic activity, as already noted, will be considered epistemic “betting” - a doing with an unpredictable outcome (Muldoon and Weisberg 2011).

“Normative pluralism” (discussed in section 4) will specifically designate the difference in norms guiding the search for knowledge such that it allows for the agents working “under” different sets of norms to have *different local peaks* in the epistemic landscape (Page 2008, 157). The agents and norms have to be sufficiently different from each other to allow for the division of epistemic labour and thus for certain agents (and normative communities) to make *relatively suboptimal, and thus redundant*, epistemic bets. Normative pluralism will be argued to be of fundamental epistemic value for the design of superior social epistemic systems, and thus *epistemically instrumental*.

“Social epistemic exclusion”, discussed in section 6, will be exemplified by the indefinite lack or systematic denial of a) sustenance, and b) epistemic resources and possibility for Epistemic Contribution, as relating to individual agents. In terms of very provisional indicators, satisfaction of sustenance would *minimally* comprise of access to food, water, energy, shelter and healthcare, while the satisfaction of access to epistemic resources and possibility for Epistemic Contribution would *minimally* comprise of basic liberties (freedom of speech, association and trade, voting rights, political freedoms, freedoms from arbitrary imprisonment), free access to epistemic materials (free school, free university, free internet, public library, strong open access and open source policies, and limited intellectual property), certain conditions of low-cost market entry (lack of formal denial of market entry on arbitrary grounds, free access to limited intellectual property rights for low-status contributors, antitrust laws and laws against abuse of dominance), and access to discretionary time, in order for the agents to be able to diversify their Epistemic Contribution. “Universal social epistemic inclusion” will designate the satisfaction of these indicators for all agents in a population. “Epistemic Contribution” may also be more broadly described, following Miranda Fricker, as “the exercise of (...) social epistemic capability on the part of the individual to contribute to the pool of shared epistemic materials – materials for knowledge, understanding, and very often for practical deliberation” (Fricker 2015, 76). The more developed understanding of social epistemic exclusion is an objective for institutional epistemology, and further

investigations may show that the expansion of indicators is required. For the present purposes, however, the minimal version will suffice to make the detection of social epistemic exclusion possible. Furthermore, this is a prescriptive, and not a descriptive, account of social epistemic inclusion. It will be argued that if an institutional arrangement cannot protect agents following redundant norms from social epistemic exclusion, it is epistemically defective. When assessed by the levels of satisfaction of these indicators, the real-world institutional arrangements may be recognized as failing to provide conditions conducive to the search for knowledge. I will focus on the discussion about the tool for tracking these conditions in Chapter 3.

I will now proceed to introduce Hayek's understanding of our constitutional ignorance.

1. 2. 2. Ignorance

Throughout his career as a philosopher of society, and a pioneer in institutional epistemology (Anderson 2006; Boettke 2018, 30⁵), Friedrich August Hayek maintained and elaborated two key ideas: 1) that no agent can have access to the totality of knowledge and that therefore no single agent can have a *categorically* superior position in the search for knowledge, and 2) that knowledge is dispersed, distributed unpredictably and unsurveyably within the population.

The first claim is one of *necessary suboptimality of epistemic agents* - it sets all epistemic agents on equal footing with regards to them having limited and insufficient evidence, limited and error-prone inferential capacities and a limited and imperfect conceptual apparatus for navigating the dynamic and complex epistemic landscapes (Hayek 1945). They *do*, however, have access to some evidence, use some concepts and make some judgements - and through this, they engage in the search for knowledge with *varied* success. This variation is crucial - it shows certain

⁵ Boettke specifically recognizes Hayek's project as *epistemological institutionalism*.

epistemic actions have better consequences than others. These actions are, however, not all available to any single epistemic agent, nor are they predictably distributed among specific epistemic agents which can be identified with certainty before the actions have been tested. The knowledge is then, as the second claim would have it, dispersed. The search for knowledge as a social enterprise revolves around the question of how to continuously and adaptively harvest and streamline all of these bits of knowledge into solutions to problems. It is the question of “utilization of knowledge not given to anyone in its totality” (Hayek 1945, 520).

This is one of the founding accounts of the design and assessment of social epistemic systems. The search for knowledge, in Hayek’s account, ceases to be an activity exclusive to an individual, and becomes a question of *the system*, of an institutional order which is most conducive to a superior epistemic performance of a population of epistemic agents (Boettke et al. 2013). The continuous and dynamic navigation of the totality of the epistemic landscape will demand a system which extracts knowledge from as many agents as possible, in order not to miss out on the Epistemic Contributions they may make in various unpredictable setups (O’Driscoll and Rizzo 1985, 41). It will moreover demand an understanding of the condition of ignorance under which individuals and systems search for knowledge.

For Hayek it is paramount to comprehend ignorance as “*constitutional*” (Hayek 1978a, 5) of epistemic agents - including those designing and assessing a social epistemic system. Properly understood, ignorance is an “irremediable” (Hayek 1982, 12) state of epistemic agents. No agent can have access to the totality of knowledge nor to the undoubtedly best strategy in the search for knowledge - and each agent makes epistemic actions *from this state*. If this is the case, “constitutively”, as the founding feature of any agent, all epistemic actions taken by the agents are to be regarded as *bets - epistemic actions under conditions of epistemic suboptimality*. Hayek believed that the judgements must be tested and contested, and thus made within an experimental environment open to disagreement. “Man acted before he thought and did not understand before he acted.” (Hayek 1982, 18) Hayek never used the term *epistemic betting*, but this concept appears fundamental if we take the considered suboptimality of epistemic agents as given.

Ignorance cannot be expected to be reliably properly recognized when at play. Agents have no immediate access to the world outside of their epistemic world - they bet on the states of affairs and infer (and thus *bet again* on the interpretation of) consequences from the new, resultant state of affairs⁶. In Brandom's vocabulary (2001), agents make normative commitments - they take on responsibility and exchange reasons for following a norm when faced with a contest among norms. Moreover, the "world outside" includes other agents, obviously, but also includes agents themselves. Ignorance, in other words, also extends to ignorance of one's best interest, one's best bet as to the satisfaction of the interest, and one's evaluation of satisfaction of the interest. The ignorance is ultimate even when judging of one's own judgements (O'Driscoll and Rizzo 1985, 48). It is also ultimate, of course, when judging any judgment. There is no "higher plane" from which the bets are assessed. The optimality of an epistemic state cannot ever be assessed by any epistemic agent without this assessment itself being an epistemic bet (Hayek 1982, p. 5).

This last point strongly advances the notion that the designer and the assessor of the social epistemic system is, constitutively, as ignorant as any other epistemic agent. This leads Hayek towards a methodological project of urging the social sciences and governance to abandon "constructivist rationalism" (Hayek 1978b, 3-23; Hayek 1982, 5). "Constructivist rationalism" for Hayek represents any program in institutional epistemology built on the notion that the superior social (epistemic) system can be designed by "reason" (Hayek 1982, 21), a *centralized* construction and production of beliefs and norms from particular premises. Hayek's social epistemic project is fundamentally opposed to delegating the totality (or the most relevant amount) of epistemic tasks to a centralized epistemic body, a group of experts or elected officials. He recognizes this as the crucial flaw in designing a social epistemic system (Boettke et al. 2013, 11)⁷.

⁶ See Sellars 1963 for a more detailed understanding of the described pragmatist dynamics, and Aligica 2014, 78 for an account of the search for knowledge as the pursuit of adaptive isomorphism between the cognized and operative environment of the agent.

⁷ Even if they were fully benevolent and outstandingly wise, the central board of planners (experts or elected officials), to whom all epistemic labour is delegated within a centralized system, would still firstly foreclose the ability of the population to make use of dispersed information (and thus epistemically unjustifiably reduce the inferential and conceptual capacity of the social epistemic system in the search for knowledge), and secondly are still ignorant (only supposedly relatively less so according to *particular* standards) and still themselves may only bet. With regards to experts' ignorance, see Gaus 2008 and, on the particular issues with the foundational problem for development of expertise with regards to the wider concerns of policy and system design, see Rittel and Webber 1973.

Within Hayek's account of the design and assessment of social epistemic systems, however, this epistemic illegitimacy of the centralization of claims of optimality obviously does *not* imply there cannot be claims of optimality within populations of agents - there obviously can, and, more to the point, there *must* be (Hayek 1945, 520-521). Without claims of optimality no search for knowledge is possible. Moreover, despite the fact that such a large and relevant aspect of epistemic lives is conceded to ignorance, Hayek's project *is* still an account of the design and assessment of the search for knowledge in the large, complex population. Exercise of restraint from centralization and epistemically distortive control over mass epistemic betting (Aligica 2014, 50; Krstić 2012, 134-135) is among its design principles, and it is precisely the humble *institutional* overcomings of irremediable ignorance on part of any possible epistemic agent that the design and assessment of the social epistemic system must account for (Fleetwood 1995, 132; Boettke et al. 2013). But perhaps even more importantly, agents are agents because they make claims of optimality, epistemic bets, and some agents make both really and evidently better claims of optimality than others, and their claims of optimality are superadditive (Hayek 1960, 52) - in interaction and building upon each other, they may move away from suboptimality as well as towards real optimality. The judgement made by any agent considering their individual and overall success has to, however, remain understood as an action and as *an epistemic bet* - another move in the *unsurveyable* game of epistemic development. *It may be a good bet.*

Norms, rules of the various games played in a large and normatively complex population, are very relevant bets - they are instructions on behaviour within certain sets of situations, the cognitive and social habits of evidence-gathering, evidence-weighting, inference and assessment of action. These bets are of particular interest for the design and assessment of the social epistemic system because they are bets on betting *strategies*, meaning they are rules which designate the sets of epistemic bets available and favoured within particular sets of circumstances.

1. 2. 3. Norms

Literature on normative behaviours as conceptual tools to make sense of the social world has been steadily growing for the past half a century, most prominently in the development of the Hayekian project into the larger family of new institutionalisms (see North 1990; North and Denzau 1994; Ostrom 2005; Hodgson 2006; Wallis 2011; Tarko 2015) as well as in philosophy of society (see Bicchieri 2006; Bicchieri 2017; Gaus 2013; Guala and Hindriks 2015; Brandom 2011). The present account will focus on the most relevant aspects of norms on which there is a large consensus among various theoreticians and philosophers. I will use the term “norms” as the name for the family of normative behaviours en large, encompassing all the particular analytical levels of normativities and rule-guided epistemic and social action, and overlapping with Hayek’s “rules” (Hayek 2015, 232-257, 278-293; Fleetwood 1995, 106-125; Gaus 2013). I will not discuss various intricacies within numerous ongoing debates, since they are not of present concern. The key premises with regards to norms will be that 1) they are implicit and explicit rules of inference and behaviour used by agents in order to reduce their ignorance, 2) they play to the agent’s advantage, 3) they may change but fundamentally tend to be preserved, and 4) they may indefinitely preserve a suboptimal strategy in the search for knowledge.

Norms are the means by which populations of agents attempt to reduce their ignorance and the environmental unpredictability by following specific orders of action within an epistemic situation (Fleetwood 1995, 5). The “whole rationale of the phenomenon of rule-guided action,” Hayek claims to rest on the “inescapable ignorance of most of the particular circumstances which determine the effects of our actions” (Hayek 1982, 20). *Norms are why our ignorance is not fatal* (Fleetwood 1995, 132). Norms are habits of mind, guides to action, inferential practices, *betting strategies*, socially incentivized patterns of methodological directives in the search for knowledge, some of which work and some of which survive despite not working. Norms play out as regularities of behaviour (or judgement) performed by a number of agents through time, and their function is not only to reduce the agent’s own uncertainty as to how to bet, but to inform the agent as to how other agents will bet. They make other agents more predictable, and in turn allow for (a varying) higher degree of coordination as well as a more efficient transmission of information. The

largest amount of norms are considered to be used by agents “without thinking”, as tacit, implicit and non-demonstrable (Brandom 2001, 21; Rietveld 2008). Norms involved in the search for knowledge need not take form of exclusively epistemic norms or be explicitly comprehensible to the community acting under them as such - they may be a complex entanglement of moral, epistemic, aesthetic and numerous other “types” of norms, and thus hybrid in their configuration while eventually epistemic in the outcome they produce.

As agents may make errors, persistent mistakes, thus conserving their suboptimal epistemic state, they may, both cognitively and socially, follow norms which *do not* advance their search for knowledge. Some other norms, however, *do* - they show agents how to act within certain situations in order to make the best decision, solve the problem and attain knowledge.

Hayek found that following norms in the state of ignorance played well with the agents’ “calculus of advantage” (Vanberg 2006, 15), giving agents instruction on how to bet when lacking relevant information and under severe risk of error of judgement in every new situation. The introduction of norms and suboptimal agents was a relevant development in economics not because it showed that agents played against their interest but because their interest, and the way they compute their interest, had to be severely redefined by bringing into the account a set of action-outcome mappings spread through time and within groups of agents (North and Denzau 1994). In other words, where neoclassical theory saw a single agent with perfect knowledge playing out each new situation to maximize its own specific benefit, Hayek, institutional and evolutionary economics, game theory and political economy began to see individual agents with limited knowledge and capacities for attaining it, guided by rules inherited from and enforced by a group of agents which allowed for the maximization of the the agents’ benefit, as well as that of the group, in a varyingly open and unfolding set of situations (Ostrom 2000; Axelrod 1986). The image became complicated, as it did when pragmatists started to investigate the nature of knowledge (Aligica 2014, 177). The normative nature of our social and economic lives became particularly undeniable with experimental results. In the ultimatum game, for instance, the proposer can offer the responder a fair or an unfair split of some amount of money, and the responder can either accept or reject the

offer. If, however, the responder rejects the offer, they both get nothing. A utility maximiser responder would always accept. The experimental evidence shows otherwise - responders overwhelmingly rejecting an unfair split, and thus minimizing their immediate benefit (Roth et al. 1991). This suggests the responders punished the proposers for breaking the rule (in this case, of fairness). Moreover, it suggests they found the rule to be so relevant they were ready to enact punishment even at a cost to themselves (Gaus 2011, 111). The key to understanding these results was that agents' did not have perfect knowledge and therefore more often than not played "a long game" with regards to their benefits - they accrued them through time by playing by the rules as opposed to choosing the most lucrative option at every one-off situation (Vanberg 2006, 13-14). Given they were ignorant and they are to remain ignorant in every future situation, it could be argued that agents are more likely to play by the rules than choose the most lucrative option for (at least) two decisive reasons: 1) the recognition of the most lucrative option by the players may have been impossible due to the lack of relevant knowledge, and the rules may have been the only means available for distinguishing the quality of options (Heinrich 2009, 10), and 2) if all players counted on that all players played by the rules in future, they could all, reasonably, expect to have more chance at maximization of their overall quality of options (Gaus 2011, 105). These reasons however need not be a conscious, explicit deliberative action on part of any agent. They merely followed norms without any pre-normative consideration (whatever that would be). Norms are primarily implicit, and it requires epistemic effort to make them explicit. But perhaps more to the point, Hayek found agents to be "made" of norms (Hayek 1982, 36) - and thus lacking any other guides at their disposal to advance in the search for knowledge. "Individual rationality is a function of social norms." (Sunstein 1996, 909)

Epistemic agents thus constitutively rely on norms. They make the environment predictable, the coordination among agents possible, the evidence recognizable, the concepts available, and the choice of inferential practice to follow efficient. Abandoning and changing norms is therefore not a light exercise - norms give order to a messy world, and agents are invested in their preservation. The community is likely to punish the abandonment or replacement of norms, and the eventual normative innovation will take time. Given norms take place in networks and

bundles, significant adaptation of other, related, norms may be necessary to fit the new norm into the lives of a sufficient amount of community members (Bicchieri 2017, Eposito 2014). However, norms are also epistemic bets themselves, and they are and can be in certain competition with each other, being tested out by agents in every new situation, and at times being recognized as failing to produce the expected consequence. They may become useless or detrimental to epistemic enterprises *without* agents responding properly by revising them - but they also, crucially, *may* change (Kingston and Caballero 2009). Hayek furthermore recognized that they may change as a result of a conscious effort on part of the agents (Gaus 2013, 6).

They, however, fundamentally need not change. Suboptimal agents may never discover a norm is flawed, or may punish, suppress or ignore the agent claiming it to be such. *Epistemically suboptimal norms may be preserved indefinitely by epistemically suboptimal agents.* They may continue to be used by the agents despite their evident suboptimality for various reasons, but primarily because agents have no other choice - they have no other norms available by which to judge the norms they're presently using as suboptimal, no other norms with which to recognize the evidence that the consequence of their actions are failures in the search for knowledge.

The following section will argue that the only way to give the agents some slight chance at changing a suboptimal norm, and changing it consciously, *without counting on epistemic luck and inferential uniqueness of agents*, must be by *making norms which offer alternative routes through the epistemic landscape available.* If epistemic agents are both suboptimal and normative, the pluralism of norms is the minimal protection of the social epistemic system against the conservation of the most suboptimal normative strategy. The superior social epistemic system thus features normative pluralism.

1. 2. 4. Instrumental Normative Pluralism

A society which centralizes epistemic betting, and delegates it to particular agents or communities, thusly reducing the number of available and performed bets within the population, is the central subject of Hayek's criticism, because it necessarily leads to worse epistemic outcomes for that population. This is so because scarce epistemic betters (agents or communities) cannot harvest and coordinate the dispersed knowledge without resorting to the burdensome process in which knowledge gets severely distorted and its (complex) epistemic value significantly decreases. The population in which as many members are allowed to harvest and signal the harvested bits of dispersed knowledge to each other is expected to deliver epistemically superior outcomes primarily because the amount of (information-harvesting, information-processing and bet-making) epistemic agents and bets is increased, thus allowing for more efficient and epistemically more valuable feedback on the bets. The social epistemic system "learns" better if as many agents are allowed to bet.

Hayek's decentralization of betting by itself, however, does not provide a severe enough threat to the conservation of a suboptimal betting strategy, nor a robust enough protection against it, *unless* it entails normative pluralism. The key epistemically distortive feature of a centralized epistemic system is that it operates according to a single normative strategy, and thus forecloses the diversity of norms. If agents are normative and suboptimal, as Hayek recognized them, institutional epistemology should endorse the claim that the decentralization (as the granting of free access into the search for knowledge to as many agents as possible) reduces the likelihood of the population being "stuck" on the wrong path through the epistemic landscape *only if* agents operate under different norms. If agents were to operate under the same norms, there would be almost no use for their numerousness - they would act as little more than a single agent (Page 2008, 153). Normative pluralism is the minimal systemic protection from the indefinite insistence on pursuing a failed trajectory in the search for knowledge (Kitcher 1990, 17; Ostrom 2005, 44). As Gerald Gaus, himself having developed his remarkable defense of pluralism by drawing from Hayek's work, remarks, "it is through confrontation with the alien that we appreciate our presuppositions" (Gaus 2018b, 63). No principle in

institutional epistemology dealing with epistemically suboptimal agents can guarantee epistemically optimal outcome, and neither does normative pluralism (D'Agostino 2009). It does however allow for the population and the system to hedge their bets, and thus decrease the likelihood of a suboptimal "lock-in".

In order to be epistemically valuable, Hayek's decentralization must then imply the diversification of normative communities. It is the centralization of normativity, a single betting strategy available to all agents, which presents the basic threat to the search for knowledge. Normative pluralism is thus *epistemically instrumental* - it is conducive to a certain, namely epistemic, good. It is not argued as a value in itself, but as a means to an end.

An additional important clarification of normative pluralism is necessary. Institutional epistemology recognizes that epistemically instrumental normative pluralism must include *redundant* investigators. Even when an epistemic task features a complex problem with a *single* solution, normative pluralism necessary for the increase in likelihood of finding that solution must include *relatively suboptimal* normative strategies in order to allow for the contest required for the protection against the conservation of the worst strategy (Zollman 2010; Kitcher 1990; D'Agostino 2009; Ostrom 2005, 284). If eventually a community X is the one which "finds the solution", this is a result of the community X's search for the solution taking place *under the conditions of redundant normative pluralism*. But moreover, as Mill understood, for the "solution" to be "right" it must be discovered under the conditions of contest or it would otherwise be epistemically lucky and unjustified. Free contest implies redundant normative diversity - the "right solution" must withstand continuous disagreement (Mill 2003). This is the case, to repeat, when the epistemic task *does not* require "a tie" between norms from the perspective of Reason-as-such. When it does, as in numerous cases of complex and particularly wicked problems, the redundant normative pluralism retains its epistemic value obviously - it *is* the "tie".

Thus, the principle of instrumental normative pluralism as derived from constitutional epistemic suboptimality of epistemic agents states that:

the protection of a social epistemic system against the conservation of the inferior normative strategy necessitates the maintenance of normative pluralism within that social epistemic system, where the pluralism in question fundamentally includes those epistemic agent which pursue normative strategies which are *relatively more* epistemically suboptimal.

Now, pluralism of norms does present a particular problem with regards to the very rationale behind norms - namely, their use in coordination and communication among agents, and particularly their use in communicating evidence among agents. If the population consists of numerous groups of normatively separated agents, how can these groups share and add to each others' advances in the search for knowledge?

Firstly, while it may appear the separation between groups could be theoretically argued to be so strict that nothing leaks out or in, the severity of this predicament must be loosened up a bit due to several reasons. Given the complexity of normative networks under which groups and individual agents operate, they are not perfectly coherent nor finished, and agents within a normative community can and usually do work with a sufficient, and not complete, overlap of norms, allowing for *variable* receptivity to new norms among agents. Moreover, agents may belong to several normative communities, normative communities may not be in any normative conflict but merely guided through a wholly different area of epistemic landscape, and normative communities may establish cooperation among themselves in various circumstances. Finally, several normative communities, or even all within a population, may share some norms and deeply disagree with regards to others - thus, for instance, when it comes to human epistemic agents, they may share a language while starkly differing in values. However, the issue, even when loosened up properly, remains when it comes to understanding that communication of evidence should occur between normative communities vastly distant from each other - meaning their normative commitments are largely non-overlapping, their agents are largely unknown to each other, their chances of engaging on some common normative territory or recognizing the same favoured destination in the navigation of the epistemic landscape quite slim. The pluralism of norms seems to have a diminished epistemic value if it does not entail agents effectively being able to

change norms under the pressure of a better norm - and the strongly separated normative communities seem to be both epistemically valuable insofar as they pursue alternative paths and therefore allow for the more thorough clearing of the epistemic landscape, and epistemically void because neither will be able to offer the evidence they see as valuable and worthy of abandoning the alternative path.

The social epistemic system must include means of communicating epistemic content “found” by agents which have not been following the same rules and who are not known to each other. There must be some way of establishing a robust *transnormative* communication between agents. Following Anderson's trinity of intelligence harvesting systems (Anderson 2006), where Dewey saw votes and talk as the tools for such purposes within an institutional system devoted to experimentalism and thus sharing evidence and reasons on best and worst practices through weighting policy, strategy and tactical alternatives, Hayek saw *prices* as the “telecom system” (Fleetwood 1995), signals of value and amount of resource as well as value of solution that can be taken as evidence for coordination among communities without those communities sharing any highly contested or substantial rules.

1. 2. 5. Prices

Prices for Hayek play a conceptual role of signalling novel epistemic content between normatively diverse agents and groups. They allow for compressed information on the change in the state of affairs derived from particular epistemic agents to travel distances. It is relevant to understand that Hayek's account develops within a particular debate in philosophy of economics, the one in which goods and services (and with them epistemic content) may be either distributed by a central planning body or by individuals within a population. The second option, the one Hayek advocates, *can* offer communication of the dispersed knowledge, whereas in a centrally planned system signalling unique epistemic content among agents or

communities is simply implausible. In fact, in this second option, the agents are incentivized by potential gains to offer their unique information to the population, the information that is the only stuff of the social epistemic system (Sunstein 2006, 106). The unique information about a certain solution is condensed into a price signal and dispatched to the vast amount of unknown agents. This is done by every normative community, for instance, attempting to reach a particular set of destinations on the epistemic landscape (solve a particular problem). The results of the competition between normative communities are reflected in prices changing according to the quality of the solution, with plans and beliefs of agents revised in light of information they provide. In this account, prices as communicated evidence of the state of affairs direct the activities of agents, guiding them to practices recognized by the population as the most beneficial by rewarding it, and thus coordinating the vast epistemic enterprise that is the economy, with sufficiently low sensitivity to a particular normative strategy of a particular normative community.

The primary failing of prices as an intelligence harvesting mechanism, and transnormative communication devices at that, appears to rest in the notion that the competition between solutions they orchestrate need not proceed in an “orderly” fashion. Prices may lie, and the suboptimal solution may be preserved indefinitely. The competitive situation may be distorted by many more factors than only central authority - to name just a few, it may be distorted by asymmetrical power relations between communities, by inadvisably conservative impulses of the majority of communities, by mere accidents never properly recognized and fixed, by prior exclusion of the agents with key information from the game due to them making a bad bet in the past. But it is rather expected that intelligence harvest mechanisms themselves were imperfect. Given that in Hayek’s view the only alternative was central provisioning or distortion of the prices through various forms of regulations, prices lying as a result of ignorant agents behaving ignorantly, or of some unfortunate arbitrary event, may appear as the least of all evils.

Moreover, prices appear to be able to help explain only a part of the epistemic traffic that does and/or should occur among normative communities, and they appear to be embedded in the networks of *norms* (Fleetwood 1995, 60-65, 125-134) operating *between* normative communities in a way which doesn’t so much reduce their

instrumental value (since it appears unviable there will be an intelligence harvesting mechanism which is somehow not subject to certain norms however robustly it may be of use for various normative communities) as it demands further detailed investigation within the design and assessment of social epistemic systems.

None of this, however, is to be regarded as a fatal argument against prices - if it is expected that *no single* intelligence harvest mechanism can take care of all kinds of epistemic contributions.

As with votes, there is a possible reinterpretation of prices available, as shown by Sunstein (2006), by applying them to prediction markets. The price here functions more definitively, and straightforwardly, as a valuation of a bet, with the money amount signalling certainty. Putting a bet on a solution “creates a market” - other agents may bet on it as well, and thus see the rise in market price as the signal of other agents’ degree of certainty (and furthermore, the noted investment in correctness) on that particular bet. In practice, this “magic of gambling” seems to work imperfectly, but remarkably well (Mann 2007).

In sum, just as votes and talk, prices are a high rating tool for harvesting collective intelligence in a social epistemic system. The question, just as with votes, is not whether prices may or may not communicate relevant epistemic content across a normatively pluralist population, but in which circumstances and how can they do it better or worse. Robust transnormative communication will take place across a large complex network of media, and not through a single perfect channel. Intelligence harvesting must itself be diversified in order for the optimal (or least suboptimal) performance to be achieved.

Prices for Hayek, however, stood opposed to the central plan, and market was the institutional order which decentralized planning, allowing for the communication through prices to take place. As prices are an outstanding intelligence harvest mechanism, ripe for a detailed understanding and development of their particular sorts within particular sophisticated ecologies of other intelligence harvest mechanisms and further diverse institutional environment, markets are an outstanding institutional arrangement for adaptive feedback and division of cognitive

labour required for epistemic progress when embedded within an institutional diversity which constricts its possible anti-epistemic tendencies. As prices cannot perform epistemically with sufficient quality as the sole intelligence harvesting mechanism, neither can markets on their own offer the institutional arrangement most conducive to epistemic good. I will claim this is so because they lack an institutional mechanism to prevent social epistemic exclusion upon suboptimal epistemic performance, the institutional protection of the more ignorant.

1. 2. 6. Markets and Social Epistemic Exclusion

Hayek is regularly maligned because of severe misconceptions and superficial knowledge of his work (Boettke 2018, 1-15). He was most certainly much more sophisticated in his actual writings, and in particular his social epistemological works, than his political and public intellectual reputation would portray him, when it came to the endorsement, and the understanding, of markets. When opposed to “state” as a centralized decision-making, problem solving and planning body (Aligica 2014, 173), “markets” represent a pure institutional type defined by massive decentralized many-minds epistemic game of signalling change in circumstances (Fleetwood 1995, 129-130), allowing in particular for a more efficient upgrade of the epistemic state of the population upon the event of new knowledge.

Hayek was qualified with regards to the endorsement of a particular superior social epistemic system, and to the endorsement of markets. They are merely the least of all evils (Hayek 1960, 54). Correctly understood, it appears, market for Hayek represented an institutional order of decentralized planning allowing for the spontaneous epistemic development, as opposed to a centrally planned institutional order as prescribed by constructivist rationalism. Hayek found markets to be an institutional arrangement conducive to epistemic development due to their capacity for the discovery, preservation, transmission and coordination of knowledge (Caldwell and Reiss 2006, 364). Markets must allow free entrepreneurial entry (Kirzner 1997, 76), and this possibility to place epistemic bets, and make Epistemic

Contribution, is the crucial feature of the superior social epistemic system. Hayek did not regard normative commitments which were spontaneously developed as necessarily optimal (Caldwell and Reiss 2006, 361) - quite the opposite, they would *usually be suboptimal*, but if decentralization and pluralism were nurtured they allowed for incremental progress which would be impossible within monocentric arrangements lacking competition and diverse lines of inquiry. Within this environment of free spontaneous pluralist epistemic betting, it is *more likely* the resulting normative commitments are epistemically valuable - it is however by no means guaranteed, and the result may still be suboptimal.

Markets may be distorted by numerous other factors beside the state (Rothstein 2011, 207-227). They may solidify arrangements deeply inadequate for dealing with complex or wicked problems and indefinitely resistant to revision of epistemically suboptimal normative strategies. Institutions which need not take form of "state" as presently theoretically understood, or for that matter currently evident in the real world, may be market bodies or a result of market bodies, effectively enforcing their "constructivist" plan and *centralizing* the Epistemic Contribution. Hayek recognized the need for markets to take place within an appropriate institutional setting. As Caldwell and Reiss write, "(h)e claimed that when a system of free markets exists within a democratic polity under the rule of law, with strong constitutional protection of a private sphere of individual activity, and well-defined, protected and exchangeable property rights, individuals will have both the incentives and the opportunity to correct errors and to make the best use of the knowledge available to them, all the while preserving individual liberty." (Caldwell and Reiss 2006, 363.) Hayek was, furthermore, aware luck may lead to various levels and kinds of social and epistemic inequalities which no institutional order could register as satisfactory (Hayek 1982, 94) - and thus that markets within certain circumstances need not perform the sole function of arbiter of destinies (Hayek 1982, 87). Certain forms of minimal social security are allowed by Hayek (Hayek 1960, 405-430) but *not essential*, and moreover, not explicitly essential to the *epistemic* performance of the population. Hayek recognized markets have problems - he however never fully accounted for their epistemic defectiveness due to their inability to constrain the reduction of normative pluralism.

The key flaw of the market as an exhaustive description of the epistemically superior institutional arrangement governing a large and normatively complex population is that it fails to account for the institutional mechanism which would prevent depriving the agents of sustenance, access to epistemic resource and the possibility of Epistemic Contribution due to their relatively suboptimal and redundant Epistemic Contribution, or in other words, social epistemic exclusion upon bad betting. The institutionalization of pluralism can be achieved primarily by the protection of redundant investigators. If redundant normative pluralism is the minimal principle of design of a superior social epistemic system, its primary task is protection of material and political conditions for Epistemic Contribution by all investigators, including the redundant ones. Redundant normative pluralism is thus the function of universal social epistemic inclusion. While Hayek's favoured social epistemic system satisfies some indicators of universal inclusion, it fails to account for the provision of epistemic resources and sustenance, which in turn undermines the markets' conduciveness to redundant pluralism and thus their chance to attain knowledge.

Firstly, however, the necessity of social epistemic *inequalities* within a population must be accounted for, and its distinction from the epistemically *distortive* social epistemic *exclusion* outlined. Markets must fail (Kirzner 1997, 79) and the bad bet must be registered as such by the system - otherwise adaptation cannot take place. "It is one of the chief tasks of competition to show which plans are false." (Hayek 1982, 117) Markets need not self-correct when discovery is made (Kirzner 1997, 79), but they represent an institutional order *open to* revision of a suboptimal epistemic state because they allow for the discovery of *error in judgement*. In order for the correction to take place, an error must be *recognized* to have been made. It may be recognized only through the competition of plans - the centralized system of planning has no systemic incentive to recognize the error at all. Epistemic agents are constitutionally error-prone and they need not recognize error. They are not even likely to. However, only if they bet massively, they stand *the least* bit of chance of recognizing and communicating an error. And when an error gets recognized is when a bad epistemic bet gets registered as such. Bad epistemic betting thus creates *epistemic* inequalities among epistemic agents. Certain agents have less or more credence and legitimacy attached to their Epistemic Contribution due to their successful or less successful betting. More importantly, certain norms are more

epistemically reliable. Agents are more generally caught up in the *real*, complex and *epistemically suboptimal*, social world that allows and denies them opportunities to contribute to the epistemic game. The population is at all times at a disequilibrium of socio-economical and ecological conditions, and this certainly entails various contingent and unsurveyable systemic inadequacies. Nature is unpredictable and the agents are epistemically suboptimal (Ostrom 2005, 49). This is understood as the baseline condition of any viable design and assessment of the social epistemic system. Within these unfortunate conditions, agents may find themselves in any number of positions of social epistemic inequality. But even within quite epistemically *favourable* circumstances, epistemic inequality might as well be an epistemically *justified* development (Novak 2018).

If the “global” level of the epistemic space refers to the set of all possible problems, and the “local” level refers to particular problems, the epistemic value of social epistemic inequalities clearly shows that pluralism may be locally transient (Zollman 2010). Redundant normative pluralism argued for here is second-order, and it requires a population of diverse normative communities. In order for normative communities to form, they must retain certain discretionary rights to exclude agents in the course of enforcing their normative commitments. Without exclusion and transience of pluralism at the local level of normative communities, redundant pluralism at the global level would be impossible - normative communities could not form and social epistemic inequalities reflective of successful or reliable betting strategies could not emerge.

However, if an exclusion from a normative community entails the exclusion from the system in terms of agents being denied certain baseline conditions for participation in the epistemic life of a population in any relevant manner, redundant normative pluralism is again violated. The minimal baseline conditions of social epistemic inclusion at the level of the system are access to sustenance, epistemic resources and possibility of Epistemic Contribution. Agents deprived of these conditions are so severely disadvantaged that they can be regarded as excluded at the level of the system from the collective search for knowledge, and thus subject to social epistemic exclusion.

Two features are then requisite for a social epistemic system to protect redundant normative pluralism: 1) bad betters must be protected from system-level exclusion (globally sustained pluralism); 2) normative communities must retain discretion of exclusion - they may enforce their norms so that epistemic hierarchies, expert communities and social epistemic inequalities may form (possibility of locally transient pluralism). While social epistemic inequalities may be a collective epistemic virtue, social epistemic exclusion is a collective epistemic vice. If an institutional arrangement fails to prevent redundant investigators becoming subject to deprivation of sustenance, epistemic resources or possibility of Epistemic Contribution, it violates epistemically instrumental redundant normative pluralism.

When they discontinue the normative alternatives upon correction, markets *may* effectively exclude the agents from the epistemic game altogether unless there exists an extra-market institutional mechanism preventing this. Upon the spontaneous valuation of epistemic inputs within markets, redundant investigators *may* be denied access to sustenance, epistemic resources and the possibility of Epistemic Contribution, and thus become socially epistemically excluded⁸. Markets feature no institutional mechanism which would constrain the social epistemic exclusion of redundant investigators. The lack of institutional protection of epistemic agents from social epistemic exclusion upon their bad betting exposes markets to severe risk of epistemically distortive reduction of normative pluralism.

It must be stressed that this argument is focused exclusively on what is epistemically valuable for systems and populations. I do not argue here that social epistemic inclusion should be guaranteed because it is a right of the individual to be included, nor because it is fair, just or morally appropriate (even though it might be). I argue that the reduction of redundant normative pluralism through social epistemic exclusion makes social epistemic systems less likely to attain knowledge. Universal inclusion is thus not argued for as a moral or political virtue, but as conducive to

⁸ An additional clarification must be made. Agents are of course capable of various epistemic activities, have various epistemic resources at their disposal and may produce significant epistemic outcomes when lacking these conditions. However, the system which fails to constrain social epistemic exclusion is, from the standpoint of institutional epistemology, a system (in which agents are) *less likely* to attain knowledge. The inconsistent and defective provision of sustenance and access to epistemic resources as well as the diminished possibility of Epistemic Contribution put agents at a significant disadvantage in their search for knowledge. Their counterparts which have these conditions satisfied are more likely (but given their suboptimality, in no way guaranteed) to attain knowledge.

collective epistemic benefit.

As shown in section 4, normative pluralism must be redundant *even when transient*, and thus bad bets are necessary for it to be epistemically valuable. Most of the bets are bad, agents are ignorant, and their norms are usually suboptimal with regards to the best course in the search for knowledge. As Hayek makes clear, "(c)ompared with the totality of knowledge which is continually utilized in the evolution of a dynamic civilization, the difference between the knowledge that the wisest and that the most ignorant individual can deliberately employ is comparatively insignificant." (Hayek 1960, 82) The population of epistemically suboptimal agents must hedge its bets by nurturing a variety of redundant investigators. Withholding sustenance and epistemic resources and foreclosing the possibility of Epistemic Contribution to agents who were instrumental in the search for knowledge by following alternative paths which were subsequently found to be relatively suboptimal is fundamentally contrary to the core instruction to the system to foster alternative paths even when a single path will eventually be found to be the correct one. This understanding alone could be sufficient to find social epistemic exclusion upon bad betting epistemically distortive.

But furthermore, social epistemic exclusion upon bad betting also disincentivizes all agents from betting, thus reducing normative pluralism. The rate of innovation in societies facing dire emergencies appears to decline because the likelihood of risk-taking required for invention on the part of the agent, or the adoption of this invention on the part of the group, decreases when the cost of failing is high (Heinrich 2009). In such circumstances, societies appear to fall back on conventional social habits and conserve existing normative strategies. Likewise, on the level of the agent, if the cost of a bad bet is social epistemic exclusion, it is to be expected that agents will become significantly less likely to make risky innovative bets and more likely to conserve a suboptimal strategy. Social epistemic exclusion upon bad betting is again epistemically distortive.

Finally, those agents who follow maverick normative strategies are of immense value for our search for knowledge because they investigate those areas of the epistemic landscape that others neglect or avoid (Weisberg and Muldoon 2009). Given their

bets may easily be regarded as bad by various and particularly socially powerful normative communities, they run a particular risk of social epistemic exclusion unless they are *institutionally* protected from it. As already noted, the epistemically valuable social epistemic inequalities must be preserved by allowing normative communities to exclude agents in the usual course of enforcing their norms, while the agent “thrown out of” all normative communities in the population must continue to have the minimal conditions satisfied. It is paramount that the social epistemic system protects the eccentrics. As Mill observed, the lack of them “marks the chief danger of the time” (Mill 2003, 140)⁹.

Markets may most certainly serve an epistemic function within an institutional order most conducive to epistemic good - they merely cannot be the *only* feature of this institutional order. The mechanism restraining social epistemic exclusion can only be an extra-market institution. Research in polycentric governance appears to be a particularly valuable resource for further inquiry with regards to the development of a robust and adaptable “ecology” of diverse institutions required for the efficient self-correction in the provision of an epistemically sufficient social minimum (Ostrom 2005; Aligica and Tarko 2013; Aligica 2014, 48-49; Aligica et al. 2019). The focus here however is exclusively on the claim that the exhaustive description of a superior social epistemic system must account for the institutional protection of redundant investigators.

1. 2. 7. Conclusion

Hayek’s fundamental contribution to institutional epistemology, the understanding of epistemic agents as suboptimal and normative, strongly supports redundant normative pluralism as the minimal principle for the design of superior social

⁹ There could be further arguments against social epistemic exclusion upon bad betting. For instance, it could be argued that the exclusion cannot ever be fully epistemologically justified due to the judging of the epistemic status of the bet itself being a bet. The presently recognized bad bets may be *falsely* recognized as such (Berg 2003, 414-415). Moreover, given all agents are epistemically suboptimal, no agent can with full justification predict any agent’s future bets from their history of betting (Sunstein 2006, 87).

epistemic systems. The protection of Epistemic Contribution by more ignorant among us is a condition of knowledge. Markets alone cannot be regarded as the institutional arrangement most conducive to epistemic good unless they are embedded in an institutional diversity which protects the redundant investigators.

1. 3. Division of Epistemic Labour and Diversity Trumps Ability: The Lessons From “Simulators”

I will lastly introduce three key contemporary works in institutional epistemology. The first two, Weisberg and Muldoon’s “Epistemic Landscapes and the Division of Cognitive Labour” (2009) and Zollman’s “The Epistemic Benefit of Transient Diversity” (2010), exemplify recent investigations in the division of epistemic labour, the foundational discussion of institutional epistemology (Mayo-Wilson et al 2011), while Page’s Diversity Trumps Ability (DTA) Theorem is a classical contemporary resource for the defence of the pluralist project in institutional epistemology and political philosophy (Anderson 2006; Landemore 2013; Gaus 2016). Both division of epistemic labour and DTA advance the understanding that the pluralism of local peaks, the epistemic contest accommodative to relatively suboptimal normative commitment, is required to produce or discover knowledge.

Furthermore, all three contributions also make use of agent-based simulations to argue their cases, which represents a new methodological development in the field (Klein et al 2018; Reijula and Kuorikoski 2020). These tools allow for testing of hypothesis on features of agents and groups of agents with regards to their systematic connection to varying qualities of epistemic output (Berg 2003) - and offer models of emergence of system-level regularities from sub-system, even agent-level, traits. The suggestive findings of simulations should be understood as akin to arguments which claim that certain features X have epistemic value Q under certain assumptions (parameters) Y.

I will now proceed to explicate the key lessons from “simulators”.

1. 3. 1. Weisberg & Muldoon, “Epistemic Landscapes and the Division of Cognitive Labour”

In Weisberg and Muldoon’s paper (2009) “epistemic landscape” is a term of a specific technical nature - it designates a distribution of peaks of epistemic significance (true *and relevant* findings) among patches of possible movement. The agents are places at zero significance areas and proceed to play out possible moves towards the maximum significance areas. These take place within cycles which allow Weisberg and Muldoon to track time it takes agents with particular directives for behaviour (thus, a kind of simple *norms*) to make epistemic progress.

Weisberg and Muldoon first investigate the epistemic progress made by individual scientists with no communication among them, and thus unable to learn from each other. These “control agents” can keep track only of their own movements in the epistemic landscape, judging the significance of their own locations and deciding on where to move next in the search of significant locations based only on the information they themselves have gathered. There is no division of cognitive labour (REF). Adding scientists in this case exhibited certain benefits, but became mostly insignificant after the number of scientists reached 30. The lack of learning opportunities, namely the lack of input from other scientists on the locations and approaches they had the chance to investigate and report on among each other, leads to, *in the best case*, highly inefficient epistemic progress - the “controls” may and do eventually find significant locations, but, save lucky instances, only after a problematically long period of time. In other cases, “controls” would converge and get stuck in situations of varying suboptimality.

Weisberg and Muldoon then introduce two strong rules for agents to learn from other agents - first is to proceed exploring the landscape with the use of evidently successful approaches and the second is to proceed by avoiding the evidently successful approaches. The first groups of scientists is named Followers, the second Mavericks. Both groups, thus, make decisions on the basis of “markers” of significance of a particular approach to the landscape which are left by agents other than themselves.

When discussing Followers, certain findings should be stressed. Weisberg and Muldoon have set up two peaks (point of maximum significance) in the landscape, and the results of Followers for finding both were relevantly lower than they were in finding either one. While Followers' chance of finding the peak of maximum significance did rise with addition of agents, Followers tended to find the single peak rapidly and then abandon any further search. Even more troublingly, they also tended to get stuck following each other around a suboptimal area of the "hill" (the area of the landscape featuring a peak of maximum significance) without eventually moving towards a peak. In the final evaluation, Followers appeared to have performed poorer than did "controls".

Mavericks, the scientists favouring unvisited patches, on the other hand, even in small groups, tended to find both peaks. In groups of 20 they always found them. Moreover, addition of agents improved the Mavericks' performance outstandingly. They were by far the most successful group.

Moving away from pure groups, Weisberg and Muldoon put to test mixed groups featuring both Followers and Mavericks. Adding even a single Maverick led to a significant improvement in performance, and adding further numbers of them followed accordingly. The results are not only due to Mavericks' own performance. They appear to help Followers get unstuck and thus make significant approaches more available to them.

Weisberg and Muldoon argue for the *mixed* groups as the epistemically superior to all others. While Mavericks tend to find and/or promptly find maximum significance peaks, when put in context of real scientific work, Followers appear valuable in *thoroughly* exploring approaches and thus articulating particular paradigms, allowing for their comparison and evaluation. This matters in science, where the bundles of norms, approaches, used by particular groups compete, and thus may themselves be regarded as objects of inquiry and assessment. Furthermore, certain norms shouldn't be abandoned, and Mavericks work by abandoning norms. *In mixed groups, however, these differences become assets.* The Followers are needed for establishing a normative community and Mavericks introduce dissent required for the revision of suboptimal normative commitments.

Weisberg and Muldoon use a set of rough traits to distinguish their agents and their landscapes are simple, and they themselves alert to the possibilities of complicating the agents' features and the landscapes becoming variously rugged. However, the central strength of their claims appears unshaken by these considerations.

1. 3. 2. Zollman, "The Epistemic Benefit of Transient Diversity"

The key insight of the division of cognitive (or epistemic) labour is that there is a discrepancy between individual epistemic virtues and collective epistemic virtues, if epistemic virtues are to be taken as traits favourable to obtaining a better outcome in the search for knowledge (Weisberg 2010, 2). In an account of *transient* redundant normative pluralism, Zollman shows a conservation of a suboptimal normative commitment, an individual epistemic vice, contributes to a better collective epistemic performance. The epistemically superior pluralism thus *must* contain normative communities which abstain from revising the suboptimal epistemic state even after exposure to evidence to its suboptimality.

The situation of Zollman's concern is that of a scientific community facing two (or more) competing theories and attempting to find the proper one. They are in a state of disagreement and want to find the truth. Zollman models the situation in the following way. Scientific communities are networks of nodes among which information is shared. The nodes are players of two slot machines which each may come up "win" or "lose". The players should learn by playing which one is more likely to deliver which result. In randomized trials, the player would play both machines an equal amount of time in order to pass judgement on the machines' performance. However, this is not a profitable strategy, and agents want to optimize - their goal is to gather information on the quality of the machines but primarily to play the best machine. This situation is descriptive of the predicament numerous scientific communities find themselves in - between the choice which research line to pursue

(which machine to play) and the time needed to pursue both in order to make the best choice which to pursue (playing both machines).

Additional caveats apply as well - in Zollman's model, among other technical specificities, past successes do not influence the probability of future success, players learn through Bayesian reasoning, and they are interested only in playing the better machine. Since they are in a network within which they may gather information from other players, this last assumption entails that they wish to leave information gathering to others and be the ones to make the best bet. Among other things, Zollman notes this assumption mimics the current state of the scientific community within which scientists are rewarded for current successes, and not future.

Zollman presents three learning situations within groups - 1) the cycle, in which each node has two neighbors with which it shares information, 2) the wheel, in which an individual in the middle shares information with each individual in a cycle but they do not, and 3) the complete graph, where everybody shares information with everybody. The results show, somewhat counterintuitively, the cycle "wins" - the situation in which each player has the access to the smallest amount of information. As Zollman writes, "It would appear here that the amount of information distributed is negatively impacting the ability of a social group to converge on the correct methodology." (Zollman 2010, 34) In the complete graph, the inferior theory is too invasive - once the information on a theory is disseminated, the players settle for the suboptimal epistemic state too fast and abandon information gathering.

He then goes on to distribute the priors of varying strengths among agents, thus making those with more extreme priors less resistant to change the theory upon evidence. When the learning situations are played again, the complete graph with extreme priors performs significantly better than any other alternative. Extreme priors maintain diversity in the information-rich environment sufficiently long for the evidence to begin to accumulate. The diversity this scientific community exhibits is transient - it sticks around long enough for the players not to settle for the inferior theory, but does not foreclose convergence on the proper theory. It can thus be achieved by either limiting the amount of information among scientists or by fostering their tendency to conserve a suboptimal betting strategy.

"The offered solutions to this problem all turn on individuals being arranged in ways that make each individual look epistemically sub-optimal. The scientists do not observe all of the available information or have overly extreme priors. Looking at these scientists from the perspective of individualistic epistemology, one might be inclined to criticize the scientists' behavior. However, when viewed as a community, their behavior becomes optimal."
(Zollman 2010, 22)

The concern of the design and assessment of social epistemic systems is excavating principles for the governance of large, complex, dynamic, wicked-problem populations - and thus transiency, as will be noted in Chapter 2, cannot be regarded as the feature present in all epistemic tasks. However, Zollman shows that even when present, and the problem at hand has a single solution, redundant normative pluralism exhibits epistemic benefits.

1. 3. 3. Page, *The Difference*

Page's central claim, the Diversity Trumps Ability Theorem, was an unexpected result of agent-based modelling experiments. Significant work for the development of the Theorem as presented in *The Difference* (2008), Page's book which features the version of the argument I will focus on in the present account, was done in cooperation with Lu Hong (Hong and Page 2004). The following is the reconstruction of the account through its crucial points and leaves numerous finer details aside as discussions which, while relevant for the inquiry into the design and assessment of social epistemic systems, are not of present concern.

Page presents intelligence as a toolbox for navigation of the epistemic environment. The larger the toolbox, the more exhaustive the navigation. Groups of experts have smaller number of tools than do diverse groups. When it comes to certain problems, this small number of just the right tools is a perfect fit for solving the problem.

However, when the problems become more complex and less familiar, the rise in the number of tools overlapping with a particular diversity of tools takes center stage in the search for knowledge. The ability in these situations represents the scarcity of tools, and moreover, lacking the proper diversity, the scarcity of upgrades to tools, which are available only in the interaction of different toolboxes and are required for the unknown peak to be conquered.

Page begins by presenting the concept of the cognitive toolbox consisting of “*perspectives* (ways of representing the world), *heuristics* (techniques and tools for making improvements), *interpretations* (ways of creating categories), and *predictive models* (inferences about correlation and cause and effect).” (Page 2008, 22)

Perspectives are representations of a problem through an internal “language”, an organization of knowledge into discernible particular items, mappings of the set of possibilities in which “names of things” stand in relations to each other which render them basically understandable and manipulable.

You and me have different perspectives on how to get to the city, and we would advise a foreigner differently. Perspectives’ difference to interpretations might be stressed - perspectives require each item be given a name, while interpretations require each group of items be given a name, thus partitioning perspectives.

Heuristics are rules for “satisficing” (finding good if not best solutions [Page 2008, 58]) derived from a previous problem which was solved using that rule and directing the agent’s search for knowledge. They guide navigation within perspectives when approaching a problem, favouring the shorter route. They may be simple rules-of-thumb (Page uses George Costanza’s “do the opposite” as an example [Page 2008, 52]), but may also be complex and sophisticated strategies. Despite frequently being used indiscriminately by agents (as is the case with George Costanza), no single heuristic, Page stresses, can work on all problems.

Interpretations create categories. They are category-level conceptual structures which allow for the judgement (proposition of a solution, or of a relevant inferential step towards a solution) with partial insight into the state of the affairs by “lumping” together things according to certain features. Interpretations are partitionings of parts

of reality according to and productive of structures required for epistemic labour. Where perspectives have one word for one thing, interpretations offer one word for a group of things, and thus allow for inference of causality and connection of significantly finer scale than perspectives do.

Predictive models allow for judgement of future state of an item on the basis of its interpreted features. They enable the agent to *expect* certain environmental dynamics playing out by themselves (and in particular circumstances) and thus, certain consequences upon particular intervention into environment. The quality of predictive models, arguably, follows the quality of interpretations. As the theory of a state of affairs become more finely developed and as the tests of the theory become more exhaustive, so does the chance at successful prediction.

Taking the toolbox approach makes it possible for Page to posit intelligence of agents as a variable collection of largely upgradeable instruments for epistemic action. The upgradability of the tools does not only entail that an each agent may adopt a better tool among those available through use by other agents, but also that tools are *superadditive*, meaning they can be combined by and among agents to create novel, more useful and successful perspectives, heuristics, interpretations and predictive models.

Building upon this account of intelligence as a toolbox, Page's proper goal is to show how the diversity of the tools creates epistemic benefits. The central claim of his work is the Diversity Trumps Ability Theorem. Given four specific conditions, the groups of diverse individuals outperform groups of more individually capable individuals. These conditions are as follows: 1) the problem has to be difficult, 2) the agents have to be "smart", 3) the agents must be able to find improvements on the local optima, and 4) they have to be sufficiently diverse not to get stuck on the same optima.

The problem has to be difficult because easy problems can be solved by any able agent. The simple problems is what their specific toolboxes are attuned to solving - these are experts encountering the questions they know the answers to because knowing the answers to these questions is what makes them experts. The problems

at which diverse groups outperform experts are those that expert do not have the ready answer for - the new, difficult problems.

The agents being “smart” means that the agents don’t have an extreme amount of optima and don’t get stuck at random solutions. Problem-solvers can be defined by their peaks, and individually better problem-solvers get stuck at relatively good values and have fewer local optima. Smart agents do not have access to global optima, but neither are their peaks so numerous and scattered that they always get stuck at some arbitrary mound. They should be *still* regarded as epistemically suboptimal because they cannot find the global optimum by themselves and, crucially, because their different representation of the problem may be *relatively suboptimal* and yet fully conducive to epistemic value (Page 2008, 48). The theorem pits the diverse inabilities against a uniform ability. The agents are not devoid of tools, however, they merely don’t have the right ones.

They must however have tools which can exhibit superadditivity. The diverse group has to contain agents capable of building on each others’ solutions - they need not be able to find the global optimum (they should, obviously, not be), only an improvement on a particular local optimum. They must, in other words, have access to the revision of a suboptimal epistemic state. They cannot be wholly incapable of revising and upgrading on certain solutions.

And lastly, the agents need to be sufficiently diverse not to get stuck on the same local optimum. Again, the constraint is quite straightforward - if the group were such that the agents could only find the *same* local optima (of low epistemic quality), their toolbox could not be regarded as properly diverse. If their paths cross quickly and they remain on the same route indefinitely, the group we’re dealing with is not diverse and less able, but closer to homogenous and less able. Thus, it suffers from the same problem as do homogenous and able groups, the problem of a limited set of tools, but by definition lack that group’s ability, namely their limited set of tools is of poorer quality for the problem-solving at hand.

The homogeneous groups, those in which agents have the same tools and thus the same optima, might as well contain only a single agent (Page 2008, 153). *The*

addition of agents with the same tools doesn't lead to change in the performance of the group. This is not the case, of course, with diverse groups, where agents can explore different paths, have different optima and, crucially, may build on each other's solutions. Homogeneous groups, containing agents of same optima, do not feature this opportunity - the best each agent can do is the best the group can do.

If these four conditions hold, it follows that the group of diverse agents will outperform the group of able agents. To make it clearer, the group of agents *who are less likely to be accurate but are diversely so* will outperform the groups of agents *who are more likely to be accurate but are similarly so.*

In a critique of Page and Hong's work, Abigail Thompson (2014) eventually attempts to show that Page's "diversity" may be regarded as "randomness". It appears, however, that claiming that Randomness Trumps Ability would only strengthens their case - and thus the critique would somehow miss the mark. Thompson purports Page has set up the conditions of inquiry in such a way that the conclusion is inevitable. The group of cognitively so similar agents that they could be regarded as a single agent will *of course* be outperformed by a group of diverse but smart agents. Moreover, she claims, we might as well posit not that diversity trumps ability but that randomness does. This is backed up, she notes, by a wealth of research into the performance of algorithms in various applications. The critique might thus revolve primarily around the technical particulars of the experimental procedures used by Page and Hong, and not around the central claim or the theoretical apparatus built to account for the central claim. Advancing the claim that Randomness Trumps Ability does not weaken the pluralist project¹⁰. Thompson seems to believe both that the central claim that Diversity Trumps Ability is so trivial that it is void to argue it, and that it cannot be of any use in understanding the real-world epistemic situations. Aside from attacking it for its obviousness, however, the uselessness of the central claim to the real world remains without specific arguments. It appears that Thompson merely finds that understanding human epistemic labour in terms of normative structures akin to algorithms (for which the claim that *Randomness Trumps Ability*,

¹⁰ For a computational study on how promoting random members may improve performance of an organization in which each level features tasks significantly different from those at the previous level, see Pluchino et al 2010.

as Thompson notes, holds) is a flawed position. It remains unclear why would this be so - it does, on the other hands, appear deeply sound to investigate epistemic labour as a normative endeavour of suboptimal agents. If we take it to be such, *the groups featuring agents who are less likely to be correct are constitutively required for the search for knowledge.*

1. 3. 4. Conclusion

The aim of this investigation is to give the most general philosophical argument concerning the fundamental institutional arrangement governing the search for knowledge of a large and complex population. It purports to offer *the comprehensive, minimal and robust restatement* of the argument for the epistemically instrumental nature of redundant normative pluralism and universal inclusion derived from the historic inquiries into institutional epistemology, and thus does not confront in a direct debate any of its predecessors, but positions their contributions in an account which is offered as the foundational to the discipline of the design and assessment of social epistemic system.

Both the new developments in the division of epistemic labour and DTA show how redundant normative pluralism conditions the search for knowledge. In addition to their theoretical value, they also present a methodological development in institutional epistemology - the use of modelling and simulations as methods of gathering insights into the intricacies of interplay of features of epistemic agents and epistemic situations most conducive to thriving while ignorant in a strange world.

2. MINIMAL PRINCIPLES OF THE DESIGN AND ASSESSMENT OF SOCIAL EPISTEMIC SYSTEMS

The present inquiry is concerned with the assessment and design of superior social epistemic systems. How should we organize our societies so that we are most capable of solving urgent and complex problems? What institutional arrangement would render the populations most likely to attain knowledge?

Starting with the clarifications of the methodological constraints of any possible design of the social epistemic systems, I advance an argument that there are at least two principles of governance which make the social epistemic system more likely to be epistemically successful, and that these are redundant normative pluralism and universal inclusion. I believe the present account advances the most fundamental, robust and minimal argument for them.

The Minimal Principles Claim in institutional epistemology I will account for in this chapter is the following:

Given the Modest Epistemic Comparative Standard and Regulative Baseline Epistemic Conditions, the population which is redundantly normatively pluralist and universally inclusive (these being the Minimal Principles, following Hedge Thesis and Trivial Clause, Output Value Unpredictability Thesis and Agent Scarcity Thesis) is more likely to attain knowledge and thus necessarily epistemically superior to any other.

The plan of the text is as follows.

I will first explore methodological concerns which constrain the normative theoretical enterprise of institutional epistemology. These will consist of setting two key questions institutional epistemology must answer in order to form sound claims on the features of a superior social epistemic system. The aim of this part is to set the strong core of methodological givens from which institutional epistemology can proceed its investigations. The first is the standard according to which we judge one social epistemic system superior or inferior to another (something I will refer to as

comparative standard). I will show that attainment of truth and agent development show serious shortcomings as comparative standards, and that the most convincing contender for it is the ability of the social epistemic system to recognize and overcome suboptimal epistemic states. I will refer to it as the Modest Epistemic Comparative Standard (MECS). The second is the set of conditions within which social epistemic systems occur – more precisely, what are the features of *any possible* population of epistemic agents? I will recognize the regulative baseline epistemic conditions (RBEC) to consist of populations of finite but unknown number of suboptimal and normative agents.

I will then proceed to extract Minimal Principles for the design and assessment of a social epistemic system which are necessary (but not necessarily sufficient) for it to be superior to others, given MECS and RBEC. These are redundant normative pluralism and universal inclusion of epistemic agents. While redundant normative pluralism will rest on one thesis following from MECS and RBEC, universal inclusion will rest on three theses, first of which follows from the normative pluralism thesis trivially, and the second and the third from RBEC. Redundant normative pluralism is argued for by Hedge Thesis (The revision of a suboptimal epistemic state is more likely if we all have different epistemically suboptimal normative strategies than if we all have the same), while universal inclusion is by Trivial Clause (There will be more of us having different suboptimal normative strategies if there is more of us and we may have different suboptimal normative strategies), Output Value Unpredictability Thesis (It cannot be known with certainty in advance which agent will contribute the revision of the suboptimal epistemic state because no epistemic agent can predict with certainty the output value of any epistemic agent) and Agents Scarcity Thesis (Given that there can never be a sufficient number of agents for the revision of the suboptimal epistemic state to be guaranteed, it is in the best interest of all agents that each agent is in the best position to produce knowledge). Each thesis is discussed, clarified and qualified where necessary. Response to the central objection to Minimal Principles, what I will call the Joker Objection (Particular agents may reduce the quality of the epistemic output), is given. The aim of this chapter is to form a strong core set of fundamental claims from which institutional epistemology can continue its investigations.

I will now proceed to argue the revision of a suboptimal epistemic state as the best candidate for the comparative standard for assessing the social epistemic systems.

2. 1. Comparative Standard in Institutional Epistemology¹¹

2. 1. 1. Introduction

Institutional epistemology is the study of the epistemic performance of *social epistemic systems*, institutional arrangements governing over large and complex populations of epistemic agents¹². One of the foundational concerns of institutional epistemology is according to which property of the social epistemic system must we judge its epistemic performance - the question, then, of comparative standard. I will argue the ability of the social epistemic system to revise suboptimal epistemic states should be regarded as a comparative standard in institutional epistemology.

The difference between the comparative standard, a methodological device for achieving the task of being more likely to attain knowledge, and the task of attaining knowledge must be clearly delineated. I will argue that the system which is designed to be able to revise suboptimal epistemic states, as opposed to being design to nurture agent development or attain the truth, is *more likely* to succeed in search for knowledge. Both agent development and attainment of truth as comparative standards *decrease the likelihood* that the social epistemic system will reach a less suboptimal or the optimal epistemic state - a justified social normative commitment (Brandom 2001) to a true belief.

The present discussion differs from the one on “procedure-independent standard of

¹¹A version of this subchapter is published as Zubčić, ML. 2019. Comparative Standard in Institutional Epistemology. *Philosophy of Society* 30/3: 418-430.

¹² The population is comprised of individual agents, and presumably, of communities of agents which may be understood as a single agent (Page 2008. For authoritative work on social ontology, see List and Pettit 2011; Tuomela 2013; Gilbert 2014). For the purposes of this text, I will however use the terms “agent” for individuals and “community” for groups.

correct decision”, a concept of frequent attention in literature on epistemic democracy (Peter 2016), inasmuch as the question of concern is not whether the decision made through the democratic procedure should be judged according to some such standard or it is epistemically and politically justified by the procedure itself. The interest here lies in a broader inquiry in social epistemology - according to which standard should we design and assess the epistemic output of any large and normatively complex population governed by a any institutional arrangement? While epistemic democrats will feature prominently in this area of social epistemology, the new methodological concept was needed to distance us from the particular debates in epistemic democracy, and to allow us a viewpoint from which we will be able to judge the total epistemic merit of any social epistemic system.

The plan of the text is the following. First, the two most relevant candidates for the comparative standard of social epistemic systems, agent development and attainment of truth, will be presented and it will be shown how they fail to escape the threat of suboptimal epistemic lock-in. While agent development will be supported by the work of Robert Talisse, namely his epistemic capability perfectionism, attainment of truth will be discussed in relation to the work of David Estlund and his political and epistemological treatment of the claim that those who are more likely to attain truth should exercise political authority over others. Secondly, the ability of the system to revise suboptimal epistemic states will be derived from the objections to both agent development and attainment of truth as the third candidate for the comparative standard in institutional epistemology. By focusing primarily on the threat the first two are unable to systematically stave off, the third candidate effectively tracks how conducive the social epistemic system is to learning. It will also be shown that the ability of the system to revise suboptimal epistemic states as a comparative standard can be recognized as supported by work in pragmatism and political economy, as well as in line with recent developments of “negative approach” to institutional design as argued for by Miranda Fricker.

2. 1. 2. The Assessment of Comparative Standards: Agent Development and Attainment of Truth

2 .1. 2. 1. Agent Development as the Comparative Standard

The design of a social epistemic system based on agent development as comparative standard would posit that the superior social epistemic system is the one which allows for the best epistemic development of its *individual* agents. This may include, for instance, development of individual epistemic virtues or capabilities. I will present the case for agent development as a comparative standard through a specific argument for epistemic capabilities perfectionism featured in Talisse's *Democracy and Moral Conflict*. Two objections from the division of epistemic labour to agent development as the comparative standard will be presented – first, that there can be a combination of "good" and "bad" individual epistemic traits and behaviours which combined give a collectively better epistemic output than exclusively a combination of "good" individual epistemic traits and behaviours; and second, that there can be epistemic traits which contribute to the development of the agent but are unknown at the point of assessment, and which therefore cannot be accounted for by the assessment. Both objections point to the central threat of a suboptimal epistemic state lock-in which social epistemic system designed with agent development as a comparative standard cannot avert. While I will focus on a specific account for illustration and clarity, the objections presented can be used to argue against any design of the social epistemic system based on this particular comparative standard.

2. 1. 2. 1. 1. Epistemic Capabilities as Agent Development

I will first explicate Talisse's account of epistemic perfectionism. While his primary argument in *Democracy and Moral Conflict* is itself highly relevant, sound and elegant, particularly with regards to the discussions in the fundamental discursive

nature of epistemic agents, the focus here is on the argument for epistemic perfectionism with which Talisse is concerned in the second part of the book (Talisse, 2009, 156-192).

Talisse's primary argument in *Democracy and Moral Conflict* is that in order for individuals to develop *any* kind of epistemic life requires them to be able to exercise their capacities for reason-exchange - epistemic agents are *defined* by being able to engage in reason-exchange. This is a sound pragmatist claim. Talisse, furthermore, argues that democracy is the basic institutional arrangement which allows the agents to do so (Talisse 2009, 79-154). Once faced with the Agent Ignorance Objection which challenges the thesis that inclusive deliberation in democracy is epistemically valuable by presenting evidence of individuals in the contingent historical circumstances of particular democratic regime (namely, citizens of USA in the beginning of 21st Century) seemingly ignorant of a multitude of political and scientific facts, Talisse endorses a form of epistemic perfectionism aimed at fuller development of agents' epistemic capabilities (Talisse 2009, 156-185).

The first thing to notice is Talisse's quick concession to the argument based on the contingent historical ignorance of agents. Despite his initial argument not hinging on agents' being knowledgeable (his argument posits that without being able to engage in reason-exchange agents cannot be referred to as epistemic *at all*), Talisse does, in the second part of the book, grow concerned about agents' lack of knowledge on the subjects they are engaged in reason-exchange on. Let me, before going into Talisse's defense against Agent Ignorance Objection, first point out that from the standpoint of institutional epistemology the objection of agents' ignorance need be conceded as relevant. Epistemic agents should be taken to be "constitutionally" (Hayek, 1978a, 5) and "irremediably" (Hayek 1982, 12) epistemically suboptimal. They would be *more ignorant* if they did not engaged in reason-exchange. They are less ignorant if they have a chance to engage in the social epistemic system. They, however, remain ignorant either way. As Talisse himself recognizes, this is why there is a need for the *social* epistemic system in the first place¹³. Individual epistemic

¹³ Talisse himself claims: "Our epistemic dependence is unavoidable because each individual has limited cognitive resources. Individually, we simply cannot inquire into every matter that is relevant to our beliefs; we must at some point rely on the epistemic capabilities of others, we must *defer*." (Talisse 2009, 141) He,

agents have severely limited epistemic capacities, including the ability to recognize relevant evidence, make use of relevant data and concepts, and develop good inferential practices. Social epistemic system, where populations of agents engage in epistemic activity, as opposed to particular individual agents, is less so.

Talisse's reply to Agent Ignorance Objection is as follows. Given Sunstein's valuable epistemological insights on inferior epistemic output of isolated normative communities (Sunstein 2009), Talisse first diagnoses the epistemic life of agents in question as lacking in trans-normative interaction. Their exposure to reasons and evidence beside those their communities provide is too low. Normative pluralism is the feature of a superior epistemic system because it reduces agent ignorance. This is a sound design understanding and a sufficient answer to the objection. Talisse, however, proceeds to argue for epistemic perfectionism aimed at agent development, and posits a list of capabilities (Talisse 2009, 173-177) the "state" should foster in the agents.

Capabilities need not be defined substantively, in the manner Talisse advances. Fricker's Epistemic Contribution (Fricker 2015) is also conceptualized as a capability, however, its realization on the part of the individual is conditioned on the mitigation of social obstacles to epistemic giving. Thus, as opposed to Talisse's capabilities as substantial individual epistemic development according to a list of virtues the individual must satisfy (which are the target of this section), Fricker pursues a negative understanding of capability, one that focuses on systemic and social conditions which allow the agent to engage in the search for knowledge. The difference between these two approaches is critical. If capabilities are understood in Fricker's sense, the account of their nurture is the account of system-level properties which hinder their progress and thus need to be reformed. If capabilities are understood in Talisse's sense, the account of the nurture is the account of the normative and behavioural change at the level of the individual in order for that individual to meet certain criteria of a good inquirer. I am here concerned with the latter, capabilities as "agent development".

furthermore, rightly points out that "(...) each of us epistemically depends on an entire *social epistemic system*." (Talisse 2009, 142) He proceeds to argue for the epistemic perfectionism nevertheless.

In the next section I will argue that designing a social epistemic system according to *any* list of individually virtuous epistemic features is a flawed strategy. I will not argue against Talisse's capabilities themselves. Some of Talisse's demands are straightforward and can be justified through other reasons without recourse to his epistemic perfectionism or agent development (for instance, the availability of information) - however, all are formulated as demands *from the individual*, as "certain cognitive and dispositional norms related to one's epistemic character" which the individual need to adopt in order to properly "engage in social processes of reason exchange" (Talisse 2009, 175). Their content is beside the point here - the focus on individual epistemic development itself runs the problematic risk of suboptimal epistemic lock-in.

2. 1. 2. 1. 2. Two Objections from the Division of Epistemic Labour

Two objections from the division of epistemic labour to *any* list of individual epistemic virtues or capabilities which arises from understanding agent development as the standard according to which we should design and assess a social epistemic system are:

1. There can be a set of practices or traits which cannot be understood as "good" epistemic practices or traits at the level of the individual, but which contribute to the superior epistemic output of the social epistemic system.
2. There can be a set of practices or traits unknown to any agent (including the assessor) at t_1 which contribute to the superior output of the social epistemic system at t_2 .

The first objection may be illustrated by Zollman's work in agent-based simulations in social epistemology. In "The Epistemic Benefit of Transient Diversity" (Zollman 2010), Zollman shows how the population in which information is shared among agents with extreme priors (and thus who conserve strategies or theories despite having evidence to the contrary available) is significantly epistemically superior to the one in which the agents lack such (non-virtuous) individual epistemic traits.

The second objection is more general, and requires only the concession that at any time of assessment there is a possibility of unknown individual traits which could contribute to better epistemic output. Therefore, a population with the set of traits T and an additional trait n, unknown at the time of assessment (which represents *any* time of assessment), may epistemically outperform a population with the set of traits T. Yet, our conceptual apparatus allows only for the assessment of the population with the set of traits T to be recognized as a superior social epistemic system, while in the same time the population with the set of traits T is suboptimal in relation to the population with the T+n set of traits.

What both of these objections present are cases of suboptimal epistemic state lock-in, and they could be understood as a twofold form of a single central objection, namely the objection that *agent development as a comparative standard runs the problematic risk of suboptimal epistemic state lock-in*. As Mayo-Wilson, Zollman and Danks (2011) observe, the divergence of prescriptions for superior individual and group epistemic performance, seminally argued by Kitcher (1990), is among the founding insights of social epistemology. There can be a configuration of individual epistemic practices which cannot be described as "good" from agent-level perspective which produces better epistemic output than does the configuration of exclusively "good" epistemic practices. Furthermore, there can be agent-level traits beneficial to the epistemic development of the population which are unknown at the time of design or assessment - and, as both Elinor Ostrom (Ostrom and Hess 2007, 68) and Friedrich Hayek (1960, 414) have noted, judging the presently best state of knowledge as the standard risks suppressing the optimal development.

2. 1. 2. 2. Attainment of Truth as the Comparative Standard

The other comparative standard featured prominently in the social epistemological debate is the attainment of truth. Epistemic democrats in particular have a tendency to describe features of their favoured social epistemic system as "truth-conducive" or

"truth-tracking" (Gaus 2011, 273). Two objections to attainment of truth as comparative standard are: a) it is either conceptually empty without additional specification of the comparative standard according to which we ought to find the truth or knowledge tracked, which in turn is the controversy presently under investigation; b) or the requirement following from attainment of truth as the comparative standard is that we delegate epistemic labour to those agents who are most likely to attain the truth. I will focus on the second objection, and provide a discussion on David Estlund's work related to the question of authority of those who are more likely to attain the truth. I will argue that his arguments against epistocracy are not satisfactory, and that a stronger epistemological argument against expert governance, and thus against attainment of truth as the comparative standard is required and possible.

The objection from division of labour to delegation of epistemic labour to experts is that it, again, exposes the social epistemic system to risk of suboptimal epistemic lock-in. The pluralism required for superior epistemic performance of a large and complex population must be *redundant*, and thus there *are* agents who are more likely to attain the truth. By delegating the totality of epistemic tasks to experts, the social epistemic system is lacking means of contesting the epistemic state the experts have attained. While they are *more likely* to attain the truth, the experts will not *necessarily* attain the truth. They are, moreover, *still epistemically suboptimal*. Therefore, the epistemic state the experts attained may as well be suboptimal. The agents less likely to attain truth are denied any possibility at contesting the epistemic state due to being denied access to any epistemic labour, and thus their different and disagreeing normative strategies cannot offer any contribution to breaking the consensus. The system build on attainment of truth as comparative standard has no means of *contesting its optimality*¹⁴.

¹⁴ Truth is a controversial concept (particularly, ofcourse, within epistemology). So controversial, in fact, that it appears ill-advised to use it for the robust design of a social epistemic system. Another objection to attainment of truth as comparative standard, therefore, could be that it would lead towards too much controversy as to the nature of this particular concept, and therefore the assessment could not even begin. However, it could be argued that the design of the social epistemic system need not proceed according to any particular controversial theory of truth, but, following Estlund's "minimal" conception of truth (Estlund 2008, 25), merely posit that the best social epistemic system is the one which produces the claim "X is f" when X is f. The systematic approach to institutional epistemology, thus, need not deal with truth in the manner the first objection implies. It may merely posit truth-conduciveness (very roughly, the ability to produce "X is f" when X is f) as a formal feature of a certain procedure. I will concede this point. Attainment of truth as the comparative standard need not be defined substantially as to settle the discussions with regards to theories of truth. The

2. 1. 2. 2. 1. Strong Political and Weak Epistemological Objection to Epistocracy

The authoritative argument in institutional epistemology against the delegation of epistemic labour to those agents who are more likely to attain truth is Estlund's objection to epistocracy not being "generally acceptable in the way that political legitimacy requires." (Estlund 2008, 7; Estlund 2003, 58) I will not presently engage with the majority of the particularities of his complex and sophisticated work, but will solely focus on the aspects relevant for this inquiry.

Estlund's account is focused on endorsing what he calls the Truth Tenet, claiming that "there are true (at least in the minimal sense) procedure-independent normative standards by which political decisions ought to be judged" (Estlund 2008: 30), and Knowledge Tenet, claiming that "some (relatively few) people know those normative standards better than others" (Estlund 2008, 30), while rejecting what he calls Authority Tenet, a claim that "(t)he normative political knowledge of those who know better is a warrant for their having political authority over others" (Estlund 2008, 30). Estlund rejects the Authority Tenet on political grounds, and argues democracy is epistemologically justified because it is "better than random and is epistemically the best among those that are generally acceptable in the way that political legitimacy requires." (Estlund 2008, 7) Democracy is the best social epistemic system because it is most likely *to attain the truth* among those social epistemic systems which can have political legitimacy.

Estlund is concerned exclusively with political and moral epistemic materials – whereas I am concerned with epistemic materials in general. More importantly, Estlund's argument is concerned primarily with political authority, and only secondarily with epistemological value. He proceeds to argue against the Authority Tenet on political grounds – Authority Tenet cannot hold, because the authority of "those who know better" cannot be held politically legitimate.

problem I will focus on is that the attainment of truth as comparative standard leads to defining the substantial agent trait of being more likely to attain truth as the definitive reason to delegate the totality of epistemic tasks to those agents that feature this trait.

Estlund's larger understanding of attainment of truth as the institutional epistemic comparative standard allows for expertism to reign supreme in institutional epistemic labour for epistemological reasons albeit forbidding it for political. In "Why Not Epistocracy?", for instance, Estlund appreciates "Millian" scholacracy epistemologically but finds it politically problematic, and effectively concedes general epistemological labour to experts and retains the political-epistemological labour as a domain of the democratic. From this, it follows there is an ought *simpliciter* (Case 2016) which is known to few and should be followed in the design of a social epistemic system which lacks *political* decision-making.

Estlund does make a particular epistemological objection to Authority Tenet - Demographic Objection (Estlund 2008, 215-219) - which states that contingent groups of experts may "carry" epistemic vices or suboptimal traits which override their relative epistemic superiority to other agents in the population. This is the case - however, not for any contingent reason of suboptimal individual epistemic traits, but for the necessity of a less likely revision of suboptimal epistemic state in cases of normative centralization. Estlund's objection is too weak and, moreover, cannot withstand the *philosophical* definition of experts as those who are more likely to attain the truth, and thus remain relatively epistemically superior despite any suboptimalities they may "carry". From the standpoint of the division of epistemic labour both of these claims can be accounted for. It is the nature of knowledge that it is conditioned on redundant normative pluralism - experts themselves can attain the truth if and only if operating under the conditions of a redundant normative pluralism.

The Authority Tenet may then be rejected on epistemological grounds. There may be an ought *simpliciter* and it may be known to the few, but they cannot *know* it without epistemic input from the redundant diversity of inquirers. Thus comparing the social epistemic systems according to the likelihood of knowing ought *simpliciter* is at best uninformative in design due to its trivial claim of there being an ought *simpliciter* and some agents being more likely to know it, and at worst epistemically distortive if improperly interpreted to have no epistemic reasons why not to delegate epistemic labour exclusively to experts. I will now present this objection to attainment of truth as a comparative standard in some detail.

2. 1. 2. 2. 2. Epistemological Objection to Attainment of Truth as Comparative Standard

Attainment of truth, or tracking of truth, is obviously an epistemic task of primary importance. Knowledge is at least a justified normative commitment to a true belief. However, this claim does not necessarily translate into the attainment of truth being the proper comparative standard of social epistemic systems - it would leave us with an uninformative or confusing standard. The inquiry on "Which of these two systems are closer to truth?" would merely return us to the original question of "How do you compare which is closer to the truth?" We are concerned with the *epistemological* comparative standard, and thus the standard according to which access to the commitment to a justified *true* belief may be more available to the population. As I will argue, it will be more available to the population which is more likely to revise a suboptimal epistemic state.

The other available answers to the question "Which of these two systems are closer to truth?" could be "Which system has more agents who are more likely to know the truth?" or "Which system is run by those who are more likely to know the truth?" Thus, expert-governed social epistemic system may result from the comparative standard of attainment of truth, if we were not to understand it trivially.

There is a wealth of empirical evidence (Hastie and Dawes 2001; Kahneman 2011; Gaus 2008) that experts can and do tend towards suboptimal epistemic state lock-ins. However, the case here involves experts understood much more stringently as those agents who are more likely to attain the truth at the assessment point t_1 . Therefore, the claim is not that experts tend to get stuck at suboptimal epistemic states, but that the social epistemic system which delegates the epistemic labour to experts has no institutional mechanism to prevent them from arriving at and retaining indefinitely the suboptimal epistemic state.

The epistemological problem with delegation of epistemic labour to those more likely to attain the truth is as follows. Note that "being *more likely* to attain truth at

assessment point t1" does not translate into "*necessarily* attaining the truth at t1". Therefore, the experts will not necessarily attain the optimal epistemic state – they might attain and indefinitely conserve a suboptimal epistemic state. Having delegated the totality of epistemic labour to experts, however, the social epistemic system has no means of contesting *whatever* epistemic state the experts have attained.

Experts could be understood more or less as one really smart person, both in its skillful excellence *and* in its cognitive limitations – namely, the lack of conceptual and computational resources *if* working with the same normative strategy (Page 2008). Adding agents to the expert community may improve the speed of computation and introduce some cognitive diversity (as exists within any given population of agents [Landemore 2012b]), but will not prevent it from getting stuck at the local optimum - save adding *normatively different* and thus *non-expert* agents. Normative pluralism is a condition of the discovery of the ought *simpliciter*.

If the social epistemic system delegates the epistemic task to those epistemic agents who are more likely to attain the truth, it denies itself any *systematic* ability to recognize (and revise) suboptimal epistemic state, and thus denies itself the *hedging* mechanism against such a state - if "hedging" is understood as minimization of risk of a bad epistemic "bet" on a particular strategy for overcoming the suboptimal epistemic state.

The superior social epistemic system, ofcourse, still needs and should welcome experts – just as it needs and should welcome really smart people. They bring *individually and relatively* superior (but *system-level suboptimal*) epistemic material into the epistemic pool. The division of epistemic labour in the superior social epistemic system *does not deny expertize nor does it deny the possibility of hierarchical relations within which experts hold higher social and epistemic "positions"* – it merely does not fully reduce the epistemic labour necessary for the superior social system to epistemic labour done exclusively by experts.

The central objection to the attainment of truth as comparative standard is thus the same as the one to agent development – it runs the problematic risk of suboptimal

epistemic state lock-in.

2. 1. 2. 3. Revision of a Suboptimal Epistemic State as the Comparative Standard

In the analysis so far the suboptimal epistemic lock-in, the inability of populations of epistemic agents to revise suboptimal epistemic states, has been shown to be the primary threat the social epistemic system faces. Thus, the ability of the system to revise suboptimal epistemic states appears to be the quintessential epistemically superior feature. The system which exhibits this feature is more likely to succeed in the search for knowledge. This ability of the social epistemic system to revise the suboptimal epistemic state should therefore be regarded as the comparative standard of social epistemic systems. I will call it Modest Epistemic Comparative Standard, MECS for short.

MECS is a *regulative* standard which tracks the ability of the social epistemic system to *learn* and allows for the development of minimal conditions for satisfaction of the justification criteria of knowledge. It is regulative as opposed to *positive* comparative standards of agent development and attainment of truth because it does not posit or depend on a substantial epistemic doctrine of agent traits (either in order to develop them or to identify those agents most likely to attain the truth) but designs the system able to withstand their *worst* suboptimalities. It doesn't ask how excellent are its agents - it asks whether the system can escape the deepest ignorance of its best agents. Doing so, it advances primarily a social epistemic system which is capable, in the most robust manner, of *moving away from epistemic ills*, of *upgrading its epistemic state*, and thus a system which is capable of *learning*.

MECS cannot favour any group of agents – 1) experts and agents with recognized epistemic virtues are less likely to revise the attained suboptimal epistemic state because they can be expected to form a consensus on a particular betting strategy, and subsequently a particular epistemic state, without means of evaluating and contesting that particular strategy and state, and 2) other agents are less likely by default. Therefore, they are both less likely to revise a suboptimal epistemic state

apart then they are together. Instead of favoring agents with certain properties, MECS favours *redundant normative pluralism* and disagreement as epistemically beneficial (and instrumental) developments within a population. Furthermore, redundant normative pluralism presents *minimal conditions for satisfaction of the justification criteria of knowledge* - social epistemic system featuring redundant normative pluralism opens the Epistemic Contributions of its members to *contest*, and thus makes them justifiable. As J. S. Mill observed with clarity, the ability to revise suboptimal epistemic states, “to be set right when it is wrong” (Mill 2003, 103) is the fundamental epistemic feature of epistemically suboptimal agents. For this ability to develop, and the deepest desperate epistemic state to be overcome, mere experience is insufficient and a possibility for disagreement is required (Mill 2003, 102).

The history of institutional epistemological thought thoroughly supports MECS. It is sound, as already shown, from the perspective of the division of epistemic labour - redundant normative pluralism is fully justified by the maintenance of the ability of the social epistemic system to revise suboptimal epistemic states. Hayek posits the task of competition to be to show “which plans are false” (Hayek 1982, 117), to reveal bad epistemic “bets” agents make in the search for knowledge under conditions of irremediable and constitutional ignorance. The economists working with dynamic complex normative systems in polycentric governance studies and New Institutional Economics have long considered the ability of the system to learn and adapt to be of central importance for its epistemic performance (Ostrom 2005; North 1990; for the overview of the subject of institutional change, see Kingston and Caballero 2009). Pragmatism’s key innovations in epistemology¹⁵ are positing the testing and contesting of normative commitments as baseline epistemic practices, *the experience of error as central to epistemic development* (Brandom 2001) and learning as the key feature of the superior institutional order. Finally, MECS is in line with Miranda Fricker’s negative approach in epistemic institutional design (Fricker 2015). Elaborating the ideal epistemic institutional arrangement requires focusing on threats to the social epistemic systems, the diagnosis of and solution to the points of

¹⁵ For a more comprehensive understanding of the connections between Hayek and pragmatists, see Aligica 2014.

failure the population is *prone to*. “The ideal social organism will have a well-functioning immune system, and you cannot design one of those without a proper understanding of its susceptibility to disease.” (Fricker 2015, 74)

The objection to MECS surely cannot rest on any argument against revision of the suboptimal epistemic states as a key feature of knowledge-acquisition. MECS however may be accused of being *too* modest. Revision of a suboptimal epistemic state does not imply reaching an optimal epistemic state – a suboptimal epistemic state may be revised into second suboptimal epistemic state. The response to this objection is two-fold: 1) given the agents are epistemically suboptimal, the attainment of the optimal state is never guaranteed, and there cannot be a social epistemic system which guarantees it, but 2) the revision of the *second* suboptimal epistemic state is possible only in the social epistemic system designed to be able to revise suboptimal epistemic states. If the objection would lead towards the design of the social epistemic system such that it would be compared according to its ability to reach an *optimal* state, it would merely lead towards attainment of truth as the comparative standard, and is therefore subject to the same objection.

We should judge the social epistemic systems according to their ability to revise suboptimal epistemic states, the ability to get “unstuck” from worst ignorance. The epistemically best performing population needn’t have the smartest members nor is lead by the wisest ones. It must, however, be most likely to recognize when it is wrong.

2. 1. 2. 4. Conclusion

Social epistemic system designed or assessed according to the comparative standard of agent development or attainment of truth give rise to the problematic risk suboptimal epistemic state lock-in. Lowering this risk should be regarded as the comparative standard of the institutional arrangement governing over a large and normatively complex population in its search for knowledge. The superior social

epistemic system is the one which learns best.

2. 2. Baseline Conditions of Social Epistemic Systems: The Regulative Account

Social epistemic systems are subject to three necessary constraints in design, the given features of any population of epistemic agents which limit the plausible design of the superior social epistemic system, which I will refer to as Regulative Baseline Epistemic Conditions (RBEC). These are:

1. Epistemic agents are epistemically suboptimal.
2. Epistemic agents are normative.
3. The number of epistemic agents in the population is finite and unknown.

Following RBEC, the epistemic behaviour of agents should be regarded fundamentally as *betting* (Muldoon and Weisberg 2011; Mayo-Wilson et al. 2011, p. 662)

I will now account for these three constraints.

2. 2. 1. Epistemic agents are epistemically suboptimal

Human decision-making is marred by a striking variety of errors, biases and insufficiencies (Kahneman 2011). Human epistemic labour is "purposeful (...) problem-solving, as people attempt, in a boundedly-rational way, to process incomplete information (about the environment and the strategies of other players) in a complex and changing environment " (Kingston and Caballero 2009, 21). This I take as holding. The suboptimality in question here would, however, entail a slightly more stringent understanding – suboptimality as the inability to perform all the

epistemic labour necessary to be guaranteed the totality of the revisions of suboptimal epistemic states.

Epistemic suboptimality thus defined should hold for all possible epistemic agents – be they, for example, human individuals, human-machine individuals, groups, or machines. In matters of epistemic development and superiority, then, we are dealing primarily with varying degrees of suboptimality, and therefore relative superiority. Certain agent is more or less suboptimal than the other agent. There cannot be the superior kind of an epistemic agent – there can only be an epistemic agent superior to another according to the relative degree of likelihood of attaining the truth or revising a suboptimal epistemic state. This entails, furthermore, that the key notion for the design of the social epistemic systems would be which features make the revision of a suboptimal epistemic state *more likely* – and *not* which features *guarantee* the revision.

Epistemic suboptimality understood in this manner would comprise of at least five features: 1) normative constraint (in evidence-susceptibility and inferential practices), 2) ignorance, 3) limited inferential capacity, 4) inability to predict the future, 5) conservation. I will briefly address the first four, while particular attention and more detailed explication will be provided for the last one.

Normative constraint (in evidence-susceptibility and inferential practices) entails that an agent is navigating the epistemic landscape guided by a particular set of norms which favour a certain navigating strategy and order its upgrade (or in other words, pick out which items are to be regarded as evidence, for which belief these items are evidence, and which action, new concept or norm, or prediction follows from the posited epistemic state). Ignorance entails that an agent does not hold all the relevant evidence, concepts or data to map the totality of the relevant epistemic landscape at time t_1 to have a guaranteed revision of a suboptimal epistemic state at t_2 . Limited inferential capacity entails that an agent can make use of a limited number of normative strategies for navigating the epistemic landscape due to the limited amount of inferential (cognitive and computational) resources at agent's disposal, which makes the agent incapable of surveying the totality of the relevant

epistemic landscape at time t1 to have a guaranteed revision of a suboptimal epistemic state at t2.

Therefore, the agent is epistemically suboptimal because the limited set of strategies available to the agent for surveying the epistemic landscape is *not sufficient to guarantee* that all the relevant parts of the epistemic landscape will be surveyed for the revision of suboptimal epistemic state to be guaranteed. This holds for all agents – for instance, it holds for the best and the worst epistemic agent, as well as for the agents judging who the best or the worst agent is. The best epistemic agent (for instance, the aforementioned model of an epistemically superior expert AI) would have a limited (even if high) number of navigating strategies available, and would therefore, be subject to suboptimal epistemic state lock-in the same way any expert is. The worst epistemic agent (the one with a history of poor bets, which is most likely to make a poor bet, having a reduced inferential capacity, high degree of ignorance and following a poor navigational strategy) would be the one which would have made a significantly worse bet. Both would, nevertheless, *be betting* – the outcome of their proposed epistemic revision (or any Epistemic Contribution) is unknown before the revision has been performed, and even then, the epistemic consequences of the revision are not necessarily fully surveyable. The epistemic behaviour of agents should thus be regarded fundamentally as betting. As I have already made use of the concept of "epistemic betting", so I will proceed in the rest of the account.

In an epistemic system in which agents can predict the future, the complete enterprise of institutional epistemology (and, for that matter, economics [O'Driscoll and Rizzo 1985]) is essentially void. Since our concern is with the social epistemic system in which agents with limited epistemic capacities and information make decisions and solve problems in an unpredictable and complex environment, it should be regarded as a regulative given, an ineradicable feature of epistemic agents, that they cannot predict the future. This, ofcourse, *does not* mean that they do not engage in prediction – it is among the most relevant and elementary epistemic tasks to make predictions and plans. It merely means that their predictions and plans are epistemic bets of differing quality – certain epistemic agents make

consistently better bets, certain epistemic agents make consistently worse bets, and certain epistemic agents make bets of inconsistent quality.

In all cases, the assessment of the quality of bets is itself an epistemic bet. Epistemic agents who make consistently better bets (and in the majority of cases, better bets within a particular domain of knowledge) are not categorically superior epistemic agents – they are *relatively superior* to the rest of the population. They are not another kind of an epistemic agent who can see into the future, but better inquirers with more developed and sophisticated epistemic skills and access to better information. It is highly recommendable that social epistemic system is capable of recognizing those agents which make better bets and give them a superior social epistemic position within the division of epistemic labour. However, it is highly detrimental, as shown, to delegate the totality of the epistemic labour to them. Despite differences in likelihood of making good or bad bets, all epistemic agents bet – and therefore, no valuable epistemic outcome can be regarded as guaranteed.

This, however, does not mean there are no better and worse bets, and that there are no agents with consistent history of either. It must be also made clear that even though the recognition of the quality of bets is itself a bet on part of other agents, this also does not mean that there are no objectively better or worse bets – it merely means that agents access this, as all others, aspect of reality through “the medium” of epistemic betting. There is no other way to know which bets are better or worse, and crucially what makes those bets such, but by testing and contesting them. If there is the need for an upgrade of an epistemic state (the network of the presumed good bets), the superior social epistemic system is the one most likely to undertake it. The fraught and uneasy assessment of likely qualities of bets, *the contest*, is the elementary task of any epistemic practice, and of any social epistemic system.

Finally, agents’ suboptimality also rests in their tendency to exhibit relevant resistance to changing the betting strategies, norms, of navigating the epistemic landscape, including those which consistently lead to a suboptimal epistemic state lock-in. This feature I refer to as being conservative, and it requires a more detailed elucidation.

2. 2. 1. 1. Conservation

There are two reasons for considering epistemic agents conservative to be a baseline condition which, while rendering them suboptimal, *also strongly constitutes their epistemic nature*. First, fundamentally, conservatism is instrumental in forming and maintaining a normative community, a group of agents sharing norms for navigation. Second, conservatism is instrumental in attaining epistemic value, because it allows for robust and thorough testing and contesting of norms in question.

Normative community is a group of individual epistemic agents navigating the epistemic landscape with the use of same norms¹⁶. These agents are cognitively diverse (to various degrees) and may navigate different parts of the landscape – they, however, share a bundle or a network norms of, for our purposes, sufficient similarity to be considered a normative community. One of the central formal functions of norms is that they allow agents to coordinate among themselves by making the behaviour of other agents in the normative community predictable because they share the same norms. Reduction in unpredictability of the behaviour of other agents, save a radical environmental change, significantly reduces the unpredictability of the environment. The ability to coordinate using the available norms is more existentially valuable to agents than is the revision of a suboptimal epistemic state, if the epistemic state the community is locked in is not recognized by the community as severely detrimental to their well-being, development or existence. If we take previously discussed features of epistemic suboptimality to entail the limited capacity for processing complex changes in distributions of normative commitments among other agents (in other words, keeping track of the diversifying normative behaviour within the population and therefore burdening the coordinative capacities), the existential value of stable distribution of normative commitments is to

¹⁶ Different normative communities may share the same problem (solution to which is the epistemic goal they may share). They are, therefore, not delineated according to particular problems they are solving, but according to shared norms they use in solving those problems. Pluralism entails a certain number of communities solving certain same problems using different norms.

be held as overriding the compulsion to revise the suboptimal but not recognized as existentially threatening epistemic state – for all epistemic agents. Focusing in particular on the real world of human epistemic agents, changing of norms is a costly process. Agents are invested in the normative structure of their communities, and they rely on them for more than epistemic guidance – their social position, trust in other agents, security, motivation and, often, meaning of their experiences. Social science literature designates the difficulty of changing norms due to their historical value to various agents invested in the normative structure in question as path-dependence of institutional change (Kingston and Caballero 2009, 13). The change in the environment may not be responded to by the normative community with the appropriate change in the normative strategy, because the normative strategy is rooted in a bundle of commitments, coordinative predictabilities, habits and investments which make agents highly favour conservation and, even more importantly, make them believe other agents highly favour conservation.

In addition to this fundamental reason, conservation of norms by agents is also instrumental in attaining epistemic value because it reduces "jumping to conclusions" and allows for the exploration of norms' robustness. When information is highly available to them, agents require extreme priors in order not "to discard a superior action too quickly" (Zollman 2010, 21). "Myside bias" is "a way of dividing cognitive labour" (Mercier and Sperber 2017, 221). Furthermore, agents which persistently pursue normative strategies which appear inferior nevertheless allow for the valuable investigation of neglected parts of the epistemic landscape. While a number of normative strategies are surely of low epistemic value, the formal function of conserving the normative strategy is overall of high epistemic value. Thus, the individual epistemic vice of conservation may be regarded as a collective epistemic virtue. While in numerous cases the conservative tendency may be harmful to agents' epistemic labour, it is a trait which must be considered permanent (and therefore, a baseline condition), and even valuable in a social epistemic system with particular features which can make use of it. It nevertheless makes agents *more likely* to conserve a *bad* betting strategy, and thus epistemically suboptimal.

2. 2. 2. Epistemic agents are normative

The explication of conditions of suboptimality and conservatism already featured the notion of norms as relevant for navigating the epistemic landscape, both individually and as a group. The present condition would posit that agents are introduced into the environment and proceed to treat the environment primarily as a dynamic distribution of normative commitments. More precisely, agents enter the unpredictable and unknown environment in which their ability to exhibit normative behaviour, and appropriate normative behaviour in particular, largely determines and signals to other agents their availability for (epistemic, among others) cooperation and possibility of Epistemic Contribution and development.

Norms were discussed in Chapter 1, in connection with their role in Hayekian institutional epistemology. The philosophy of norms is a rich and growing field of exceptional theoretical work. It would be difficult presently to find a philosopher or a social scientist who would not posit social life as a distinctly normative endeavour. The discussion is particularly sophisticated¹⁷, and my present explication on norms focuses exclusively on the most general and the most non-controversial features, which all or at least the authoritative majority of authors would agree on.

Agents begin building their epistemic capacities by discriminating between events in the environment, connecting those events in order of casualities, and calculating behavioural responses to those events which bring about a desired result (whatever it may be). When such connection and calculation presents itself as viable (in any manner the agent finds relevant), a norm is established and propagated through the population, or at least the community.

The norm in question could be broken down to at least three aspects: first, it endorses particular environmental signals as events of relevance, a particular connection of events as of interest for calculating behaviour, and assigns particular

¹⁷ To stress a few resources for particularly valuable discussions on the aspects of norms relevant for present purposes, see North and Denzau 1994, Brandom 2001, Ostrom 2005, Bicchieri 2006, Herrmann-Pillath 2012, and Guala and Hindriks 2015.

epistemic weight to environmental signals in relation to the particular connection of events; second, it proposes a standard for viability and addition of a norm to the bundle of norms; and third, once established and propagated through the community, it allows the agent to believe all other agents will follow a similar cognitive and behavioural path. The first aspect presents norms as tools for establishing robust habits of epistemic betting (organizing the environmental signals into concepts, making hypotheses about the environment, recognizing evidence as evidence, and making predictions); the second aspect presents norms as tools for discriminating between successful and unsuccessful management of the environment and tools of upgrade of habits of epistemic betting which are related to a particular outcome; and the third aspect presents norms as tools for reducing the unpredictability of the environment by allowing coordination among agents.

The propagation of norms, in line with the third aspect, entails that each subsequent agent entering the community builds their epistemic capacities through certain bundles or families of bundles of normative commitments enforced by other epistemic agents in the community. Not all norms an agent follows need be enforced by the community – but numerous norms are. Those norms that are socially-enforced signal to the community that agent can be regarded as capable of normative behaviour (and therefore available for cooperation and coordination), and the agent's ability to follow them (along with the agent's epistemic production and success of agent's epistemic bets, both recognized by the community as of certain value) feeds into the subsequent positioning of the agent within the organization of epistemic labour.

A couple of additional remarks should be made. Norms which are propagated through the community need not be demonstrable as norms – they may as well be, and in numerous cases are, implicit, habitual and even unknown to members of that community (Bicchieri 2006). Norms which could be understood as epistemic may comprise of a variety of "different level" norms – from fully explicated and public rules of behaviour and organization, to a subpersonal (even *if social*) non-demonstrable norms of inference, attention and relevance. In numerous cases, epistemic norms cannot be fully conceptually divorced from other norms which do not serve a direct epistemic function, such as moral norms. In other cases, moral norms are epistemic

norms in disguise. Moreover, an agent can be a member of various normative communities, groups of epistemic agents "working under" what can robustly be regarded as the same bundle of norms, at the same time – if the norms are not directly contradictory, or if the agent in question is particularly apt at living with contradictory norms. There is a variety of possible groupings of norms and normative communities (Guala and Hindriks 2015), but for purposes of this text there is no need for the delineation – the claims I make in the text should robustly count for all relevant groupings.

In sum, agents are therefore "thrown into" an environment that is most relevantly "mapped out" according to the possible "uses" of it. The repeated uses give rise to rules of what should be considered worthy of attention, what should be changed (handled, intervened into or prohibited), what can be expected, and, eventually, what else could be of use and how rule ought to change in order for the superior use to be available. In human societies, the behaviour of this kind in general is rewarded, as is the ability to make particular uses and follow particular rules. For our purposes, it will be also relevant to note that a significant part of the agents' epistemic production, and in particular the part of their epistemic production fundamental to the organization of future epistemic production, is the reproduction and production of norms.

2. 2. 3. The number of epistemic agents in the population is finite and unknown

Any population, and any social epistemic system, features a certain number of epistemic agents. It should stand as a strong methodological recommendation for the design of social epistemic systems that this number should be posited as both finite and unknown, where *finite* entails that there is no infinite number of agents in any population at any given time, and *unknown* entails that it is impossible to predict how will this number change in any population at any given time.

There being a finite number of agents in any population at any given time is a simple proposition based on reasonable expectation of environmental constraints – even the number of superior AIs doing all the epistemic work should be regarded as subject to finitude of the environment within which they operate. We might resort to thought experiments in which this environmental finitude is expanded significantly, but it would be *counterproductive* for the design of social epistemic systems to posit an infinite environment with an infinite number of epistemic agents. If we were to do it, however, the infinities in question would again be reduced to operational relations – which infinity of agents can survey which infinity of environmental reality? Therefore, it would again lead us to something akin to dealing with sets of agents and sets of environmental realities; only less methodologically apt to dealing with real-world issues of social epistemic systems.

There being an unknown number of agents in any population at any given time, on the other hand, introduces change in the number of agents as a prerequisite for the design of a *robust and an adaptable* social epistemic system. Again, the proposition is based on a reasonable expectation of environmental constraints – human agents produce offspring, migrate and might create new epistemic agents (machines) which, in turn, might as well create new epistemic agents (other machines), and the number of all of them may, moreover, decrease. But, again, the proposition is methodological much more than descriptive – the social epistemic systems designed according to a fixed number of agent cannot accommodate the unpredictable change in a dynamic environment nor, indeed, be of use for understanding *robust* principles of epistemic governance of a large and complex population.

The condition, unlike its opposite, survives counterfactual examination. If the number of agents were to be discovered at a certain point to be infinite, the system designed on finite and unknown number would have no problem adapting, precisely because it is designed on an unknown number. If the number of agents were to be discovered at a certain point to be fixed, the system designed on finite and unknown number would have no problem adapting, precisely because it is designed on a finite number. Systems designed according to an infinite or fixed number of agents are vulnerable to unpredictable constraints of finitude and change. The system designed

on the condition of finite and unknown number can adapt to both infinite and fixed number of agents.

In sum, the social epistemic system should be considered open-ended with regards to the growth or reduction in the number of agents, but at no given time is this number to be regarded as infinite. These are descriptively sound propositions, but their value to the design of social epistemic systems is primarily methodological - regulative.

2. 2. 4. Conclusion

Institutional epistemology is concerned with the design and assessment of a social epistemic system comprising of a finite but unknown number of the epistemically suboptimal and normative agents - these are the regulative baseline epistemic conditions.

What are the Minimal Principles of the superior social epistemic governance, namely the one under which a population of epistemic agents is more likely to undertake the required epistemic upgrade? Having set up two methodological constraints - *superior social epistemic systems must satisfy MECS under RBEC* - I now turn to examination of the Minimal Principles of the design and assessment of superior social epistemic systems.

2. 3. Minimal Principles

I will now proceed to present Minimal Principles for the satisfaction of MECS under RBEC in a large and normatively complex population - namely, redundant normative pluralism and universal inclusion of epistemic agents - by arguing them through four theses. These are:

Hedge Thesis (HT): The revision of a suboptimal epistemic state is more likely if we all have different epistemically suboptimal normative strategies than if we all have the same.

for redundant normative pluralism; and

Trivial Clause (TC): There will be more of us having different suboptimal normative strategies if there is more of us and we may have different suboptimal normative strategies.

Output Value Unpredictability Thesis (OVUT): It cannot be known with certainty in advance which agent will contribute the revision of the suboptimal epistemic state because no epistemic agent can predict with certainty the output value of any epistemic agent.

Agent Scarcity Thesis (AST): Given that there can never be a sufficient number of agents for the revision of the suboptimal epistemic state to be guaranteed, it is in the best interest of all agents that each agent is in the best position to produce knowledge.

for universal inclusion of epistemic agents.

I will now offer arguments and comments for each thesis.

2. 3. 1. Redundant Normative Pluralism

Hedge Thesis: The revision of a suboptimal epistemic state is more likely if we all have different epistemically suboptimal normative strategies than if we all have the same.

If the ability to revise a suboptimal epistemic state is the comparative standard of the social epistemic system in the population of finite but unknown number of suboptimal and normative epistemic agents, the best bet in the design of the social epistemic system is redundant normative pluralism. (If RBEC and MECS hold, HT follows.) This is so because a finite population of suboptimal epistemic agents following diverse normative strategies is more likely to get *unstuck* from a suboptimal epistemic state than is a finite population of suboptimal epistemic agents following a single normative strategy. More to the point, even if a certain normative community is more likely to revise a suboptimal epistemic state or follows a clearly optimal normative strategy, 1) its normative strategy has withstood the contest minimally required for a true belief to be justified and thus its discovery is conditioned on a redundant normative pluralism, and 2) it is advisable to *hedge our bets* by introducing more normative communities, namely those with *different local peaks*, into the epistemic cooperation¹⁸. The normative pluralism correctly understood must be *redundant* - the knowledge of the social epistemic system is *conditioned* on the protection and inclusion of the *relatively suboptimal, and thus redundant*, Epistemic Contribution.

In a population in which a single normative strategy is available to the epistemic agents which does not undergo a contest, epistemic agents cannot revise the epistemic state unless they are subject to considerable epistemic luck (Unger 1968; Pritchard 2005). Furthermore, those populations which contest the belief will have the belief *justified* in a *reliable* way (Goldman 1979). Normative communities are, moreover, highly conservative, and therefore reluctant to revise and inclined to postpone (indefinitely) the recognition and the appropriate evaluation of the betrayal

¹⁸ Even an individual epistemic agents would be better off in the search for knowledge if she were to test alternative theories and the best-up-to-now theory consecutively (Mayo-Wilson et al 2011, 664-665).

of expectations. This is not to be regarded purely as an epistemic vice – it may also be regarded as protecting the epistemic state against rash and misdirected revision (Zollman 2010), and as the means of establishing stable coordination, which favours the change in norms taking place in a normatively predictable environment (or, in other words, favouring the piecemeal upgrading of the network of normative commitments), given the predictability norms offer can be (not wholly unjustifiedly) regarded as more (even epistemically) beneficial than their change. However, this does not mean conservation is always a good idea – nor that it, if left uncontested, would lead to the revision of a suboptimal epistemic state. The introduction of other normative communities into the population make the revision more likely because the practices of contest of the epistemic states make the conservation of a suboptimal epistemic state more difficult. This *does not* mean conservation should (if it could) be "abolished" – it means conservation of particular normative strategies is not necessarily beneficial, and that the population is *constitutionally more likely to get unstuck from the epistemically detrimental conservation* when under conditions of redundant normative pluralism, and thus a variety of normative strategies which *contest for conservation*.

Redundant normative pluralism does not, however, guarantee the revision of a suboptimal epistemic state nor the revision of the suboptimal epistemic state for a better one. Suboptimal epistemic agents may retain the suboptimal epistemic state even under pluralism. Minimal Principles at system-level do not guarantee optimal epistemic output - they make it more likely. It is only under redundant normative pluralism that the system *might* revise the suboptimal epistemic state and thus learn. *No system design can guarantee the optimal output from a population of suboptimal agents - but only Minimal Principles make it possible* (D'Agostino 2009).

The clearer understanding of redundant normative pluralism as a Minimal Principle of design and assessment of social epistemic systems requires few more detailed inspections. First, the strong thesis of redundant normative pluralism being a minimal condition for the justification criteria for knowledge needs to be "fleshed out". It shows that hedging through disagreement, on top of protecting from suboptimal "lock-ins", is fundamental to knowing. Secondly, the relevance of interaction between normative communities for epistemically instrumental redundant pluralism needs to

be addressed. Furthermore, the distinction between globally sustained and locally transient redundant normative pluralism will be emphasized.

2. 3. 1. 1. Knowledge is Minimally Conditioned on Redundant Normative Pluralism: Justification and Reliability

Knowledge is strictly speaking impossible in a population with a non-contestable normative strategy. The free contest of the commitment and thus redundant normative pluralism allow for the minimal satisfaction of the justification criteria required for knowledge. In arguing this, I will largely focus on the key points which I find sufficiently non-controversially strong and minimal to account for the role of contest, disagreement and pluralism in the attainment of knowledge.

For a belief or a commitment to count as knowledge they must be justified. J. S. Mill's foundational insight in institutional epistemology is that a commitment must withstand contest, an *adversarial* procedure of the exchange of reasons and evidence (Goldman and Cox 1996), in order to derive justification.

“There is the greatest difference between presuming an opinion to be true, because, with every opportunity for contesting it, it has not been refuted, and assuming its truth for the purpose of not permitting its refutation. Complete liberty of contradicting and disproving our opinion, is the very condition which justifies us in assuming its truth for purposes of action; and on no other terms can a being with human faculties have any rational assurance of being right” (Mill 2003, 102).

The contest is the requirement of knowledge because it *is* a process of justification. As Brandom recognizes, the game of giving and asking for reasons is the basis of sapience, and thus any plausible epistemic activity. As Sperber and Mercier (2017) show the contest, a social act of reasoning, is what made us capable of overriding

our epistemic suboptimalities. It is *having to* justify and *be responsible* for a certain normative commitment which prompts and *constitutes* its justification.

Withstanding contest here entails disagreement in which both parties can *in principle* be persuaded¹⁹. The agents may be stubborn and not change their mind in light of better evidence and reasons, but they must be able to play the game of giving and asking for reasons. Those creatures that cannot in principle be persuaded and change their mind, and thus *cannot experience error in judgement*, do not make a normative commitment in the search for knowledge, and are not epistemic agents. In effect, without contest only a feature akin to what Hannah Ginsborg defines as “primitive normativity” of general fit between a situation and response to the situation would be available to human animals (Ginsborg 2011). If we understand epistemic action as the ability to play the game of giving and asking for reasons, and thus the ability to engage in contest, exclusive reliance on “primitive normativity” would make a creature nomic but distinctly non-epistemic. Norms which do not get explicated may play a significant role in the epistemic performance - however, it is only the norms which are explicated (as commitments) that count as an epistemic act. The non-explicated norms are not moves in the *search for* knowledge, but in *stumbling upon* knowledge. Epistemic agency is created through the explication of norms as reason-exchange. This is so no matter how low quality that explication may be. The claim that humans become epistemic agents through the explication of norms does not imply that they do it well. They may as well, and presumably do, do it badly. Does this however mean that non-human animals which do not exchange reasons are not epistemic? In a strict sense, yes. (Note, though, that it is possible we don't know how some of them do do it.) Namely, in the sense in which “justified” means at the very least withstanding contest. Non-human animals may exhibit primitive normativity, and they may even exhibit higher level normativities (Danon 2019). However, if they were to engage in the search for knowledge, it would have to be a *normative change they have disagreed about and decided to go through with*. They may emerge forms of adaptations to the environment and circumstances, they may act with *nomic and beneficial* hinge certainties (as do we, see Moyal-Sharrock 2016),

¹⁹ For a discussion on the role of being able to be persuaded by reasons in the anti-relativist reading of non-foundationalism, see Pritchard 2011.

and they may have *felt* injustice in the same way one senses red content (Sellars 1963). Only if they have dissented and engaged in a conflict in the space of reasons may they have become minimally epistemic.

Mill does not believe truth has a particular force which breaks the despotic tendencies from suppressing it (he calls such beliefs “idle sentimentality” [Mill 2003, 109]), however not only does the constitutionality of contest for understanding knowledge provide check to despotic tendencies (Kelly 2006), but the truth will always again crop up in history when some or other group of suboptimal epistemic agents “try out” the correct normative commitment while navigating the space of reasons. Save the conditions of contest, it cannot be said to be the *correct* normative commitment, if it would emerge at all.

To achieve epistemic benefit the contest here must imply the redundant normative pluralism. The quality of the collective epistemic performance is irreducible to individual epistemic virtues (Mayo-Wilson et al 2011) - and moreover, the division of epistemic labour requires a group consisting of at least an epistemically superior member and an epistemically inferior member. The paradigmatic example is the dogmatic and irrational refusal to change the norm as epistemically beneficial within information-rich populations (Zollman 2009) - the system in which the pluralism is redundant *even when* transient (in single-solution, or “tame”, problems) will exhibit the best epistemic performance. It will do so because it will allow for the contest to *play out*, and thus for the solution, or for that matter ought *simpliciter* (Case 2016), to be discovered - unless the Reason-as-such doesn't demand the tie between commitments.

Mill understood, as Hayek and Ostrom stress, that judging the present state of knowledge as optimal is the typical flaw in the design of a social epistemic system.

“To refuse a hearing to an opinion, because they are sure that it is false, is to assume that *their* certainty is the same thing as *absolute* certainty. All silencing of discussion is an assumption of infallibility” (Mill 2003, 100).

Exposure of a normative commitment to contest *allows* for it to become justified, both due to the requirements of the procedure to make reasons available and *salient (through commitment)* to all agents and to the eventual availability of reasons to all epistemic agents. Justification is conditioned on redundant normative pluralism - a contest of a *correct* commitment.

In the same vein, an epistemically *reliable* norm (method or procedure) cannot be recognized as such without withstanding the scrutiny of contest and thus disagreement. Contest produces the pressure to make reasons and epistemic procedures *explicit*, and thus *available* to other agents for testing and assessing - it allows for the reliability of the belief-forming process to be accounted for and possibly justified, and moreover to be known (Lehrer 2000). The free possibility of contest as established through the Minimal Principle of redundant normative pluralism is the elementary epistemically reliable process. The reliability of all other epistemic processes (norms, methods or procedures) rests on their discovery under conditions of disagreement. Therefore, social epistemic inequalities are justified only under conditions of universal social epistemic inclusion as the function redundant normative pluralism.

“The beliefs which we have most warrant for, have no safeguard to rest on, but a standing invitation to the whole world to prove them unfounded.” (Mill 2003, 103)

2. 3. 1. 2. Interaction

The development of interaction in pluralism is a complex task of primary importance for social epistemic systems. Without *any* interaction among different normative communities, their strategies cannot “learn” from each other. Interaction is relevant when the problem has a single solution and agents can build on each other’s distinct local peaks towards it, but also when it requires continued parallel upgrades of divergent strategies in cases of a tie from the perspective of Reason-as-such.

Epistemic value of pluralism is significantly reduced if normative communities do not interact and cannot make use of each other's advances for any purpose whatsoever. The formation of diverse normative communities entails a relevant level of parallel and isolated epistemic endeavours – but the population lacking in any zones of interaction among normative communities is seriously underserved by existent second-order pluralism, and is less likely to revise suboptimal epistemic states.

As discussed in Chapter 1, *intelligence harvest mechanisms* (IHM) are solutions to the problem of interaction among normatively diverse communities and the utilization of dispersed knowledge. They are emergent or designed social processes which enable the normatively distant or divergent communities to communicate epistemic content and effectively extract epistemic benefit from each other's searches through the epistemic landscape. The classical mechanisms for harvesting collective intelligence in institutional epistemology are *talk, votes and prices*. While both prices, as compressed messages on change signalled across considerable normative distances, and votes, as a variety of aggregations of diverse inputs, have relevant epistemic merit, the interaction as the key to epistemic development is perhaps most immediately evident in the case of talk. As delineated in the previous section, even to make their paths minimally reliable requires that the communities engage in reason-exchange with others. Talk as making normative commitments explicit, and as an attempt at translation of local peaks between normative communities as well as giving of the Epistemic Contribution across cognitive (inferential) diversity within normative communities, is foundational to epistemic agency. New forms of IHMs may also develop – for instance, prediction markets and various mechanisms for crowdsourcing intelligence appear to perform remarkably (Sunstein 2006). IHMs serving a particular epistemic purpose may require a particular design – for instance, they might focus on making room for dissenting voices (Page 2008, Sunstein 2006), or making sure the agents have some level of common understanding of the problem (Ostrom 2005, particularly 104-109), or restricting communication between agents (Sunstein 2006), or distributing weights to contributions according to the reliability (history of bets) of agents (List and Pettit 2011, 100), or developing appropriate “pidgins” between groups of researchers (Muldoon 2013). It is expected that all IHMs will be imperfect, and the diversification of intelligence harvest mechanisms in a social epistemic system is epistemically desirable.

2. 3. 1. 3. Transience

If the “global” level of the epistemic space refers to the set of all possible problems, and the “local” level refers to particular problems, redundant normative pluralism may at local levels be transient (Zollman 2010) while at the global level it should be sustained. As already discussed in Chapter 1, social epistemic inequalities brought on by local transience are a social epistemic virtue.

Cases of locally transient redundant normative pluralism show that certain problems will require a single solution, and that attaining that solution requires overcoming pluralism at the certain point of epistemic output. In other words, while we may enter the epistemic cooperation on solving the problem X with plural normative strategies, the solving of the problem will lead to “shedding” of unsuccessful strategies, particular bad bets will be recognized as bad bets and identified as useless or detrimental to attainment of the solution, and overall normative pluralism will be reduced to a single normative strategy which works (or is) best.

Even with such cases, redundant normative pluralism must be observed to be *the condition* of arriving at the best bet or attaining knowledge. However, it is to be also considered *transient*.

Transience is to be observed as a *kind* of local peaks conditioned on the redundancy holding at the global level of the social epistemic *system*. Problem X has a single solution which must be arrived at by settling for the best strategy *which withstands contest*. The *overall* bulk of problems (including those which are as yet not even conceptualized as problems, or have not even come to be), meanwhile, requires a steady influx of redundant plurality of normative strategies, in order for the best to be identified in the course of solving a particular problem. Furthermore, globally sustained pluralism allows for disagreement after local transience, which is, as Anderson notes (2006), crucial for epistemic development of the population inasmuch as it allows for the contest of and proper feedback on the best-up-to-now strategy.

“Input pluralism of normative strategies and output the best normative strategy” occurs when the superior social epistemic system (or, for that matter, normative community) is dealing with problems which have *single solutions*. Numerous complex and particularly wicked problems are not of this kind and do not require the resolution to (nor may identify) any "best" normative strategy - certain optimal epistemic states may require indefinitely competing oughts *simpliciter* (Case 2016).

2. 3. 1. 5. Conclusion

In sum, Hedge Thesis claims that the redundantly normatively pluralist population is more likely to revise a suboptimal epistemic state than is a population playing a single betting strategy. A population with a single betting strategy may be considered a single agent with (epistemically negligible) extra cognitive diversity (Page 2008), and thus cannot get unstuck from a suboptimal peak, nor, moreover, attain knowledge as a justified true belief. While locally pluralism may and must be transient, at the global level relatively epistemically suboptimal redundant investigators and their communities should be protected to develop their normative strategies with zones of interaction available, and desirable, to them.

Hedge Thesis is sufficiently strong, decisive and definitive to bear the whole weight of the argument for redundant normative pluralism. I will now move on to three arguments for universal inclusion as the Minimal Principle of the design and assessment of the superior social epistemic system.

2. 3. 2. Universal Inclusion

I will argue that universal inclusion is the Minimal Principle of the superior social epistemic systems with three thesis. The first will argue that universal inclusion

follows trivially if the redundant normative pluralism holds, the second will argue that social epistemic exclusion is always unjustified since no agent can predict with certainty the output value of any agent, and the third will argue that exclusion entails reduction in epistemic resources which are already scarce and is therefore of detriment to overall epistemic production. I will briefly explicate the first two and focus on the third, which allows for a more detailed understanding of universal inclusion. I will then present the objection to universal inclusion which claims that certain agents can be such a burden on resources that they reduce the epistemic output of other agents, and offer answers to it.

Universal inclusion entails all epistemic agents must be allowed to make Epistemic Contribution, understood as "the exercise of (...) social epistemic capability on the part of the individual to contribute to the pool of shared epistemic materials – materials for knowledge, understanding, and very often for practical deliberation" (Fricker 2015, 76). The discussion on the concept of "capabilities" will not be engaged in here – the broad understanding of inclusion in the social epistemic system at the individual agent level as having access to a variety of epistemic bets with the perspective of having access to new social betting areas according to the outcome of previous bets, will suffice for present purposes. As I will explicate further in Agent Scarcity Thesis, the key to universal inclusion as the Minimal Principle of superior social epistemic systems is for the agent *not to be denied access to epistemic betting due to a bad bet* (and in general due to a particular history of bets), while in the same time retaining the possibility of being awarded access to particular bets due to a particular history of bets. Inclusion can be defined crucially negatively as the lack of evident social epistemic exclusion, meaning wholesale denial of access to the social epistemic system in terms of agent-level deprivation of sustenance, epistemic resources and possibility of Epistemic Contribution . The particular institutional requirements for universal inclusion will be discussed in Chapter 3.

I will now go on to present the theses supporting universal inclusion of epistemic agents as a Minimal Principle of the superior social epistemic system.

2. 3. 2. 1. Trivial Clause

TC: There will be more of us having different suboptimal normative strategies if there is more of us and we may have different suboptimal normative strategies.

Building on Landemore's (2013) insight that the easiest way to increase cognitive diversity in a social epistemic system is by introducing more agents into the epistemic cooperation, redundant normative pluralism methodologically implies universal inclusion. More agents means more potential agents with diverse normative strategies. Furthermore, and even more to the point, agents are more likely to follow different strategies if they are not punished by social epistemic exclusion for following them. The true expression of pluralism is the protection of redundant investigators, including crucially those that make bad bets, from social epistemic exclusion due to their redundant strategies and bad bets. Normative pluralism is thus the *function* of universal inclusion.

2. 3. 2. 2. Output Value Unpredictability Thesis

OVUT: It cannot be known with certainty in advance which agent will contribute the revision of the suboptimal epistemic state because no epistemic agent can predict with certainty the output value of any epistemic agent.

Since epistemic agents are suboptimal and cannot know the future, as posited by RBEC, they cannot know in advance the sum of any agent's (future) contribution. The positions in the social epistemic system assigned to agents according to their histories of bets *are themselves bets*.

Neither OVUT nor any part of the present account implies that agents do not or should not arrange the significant amount of epistemic cooperation into structures which award better knowledge-producing positions to those agents with the track-record of consistently better bets. In other words, this does not mean there should not be expert structures within epistemic cooperation. There obviously should given experts are by definition those agents who are more likely to attain the epistemic value, and the existence of expert structures is fully compatible with the social epistemic system featuring a globally sustained redundant normative pluralism, and more to the point, with the social epistemic system featuring the lowest possible exclusion rate. As will now be elaborated in the course of explicating the Agent Scarcity Thesis, the institutional epistemology should demand that the epistemically non-distortive social epistemic inequalities be retained with social epistemic exclusion constrained.

2. 3. 2. 3. Agent Scarcity Thesis

AST: Given that there can never be a sufficient number of agents for the revision of the suboptimal epistemic state to be guaranteed, it is in the best interest of all agents that each agent is in the best position to produce knowledge.

Any population with a *finite* number of agents (RBEC) needs *as many agents as possible* to increase the likelihood of the revision of a suboptimal epistemic state. The particular number of agents which can be said to be sufficient with regards to the increase in the likelihood of the revision of a suboptimal epistemic state *cannot be posited*, and *neither can the cut-off point* in the number of agents at which the population becomes less likely to revise a suboptimal epistemic state. Every agent which is denied the position to produce knowledge is the potential knowledge-producing agent lost from the social epistemic system, which makes the social epistemic system less likely to revise a suboptimal epistemic state. It may be argued that the increase in the number and diversity of bets increases the likelihood of

revision of suboptimal epistemic state²⁰. However, for AST to hold it is sufficient to argue that *the decrease* in the number and diversity of epistemic bets (and, thus, agents) within populations with the finite number of suboptimal and normative agents make the revision less likely. *In populations with finite number of agents, the epistemic agents are scarce.*

As I will discuss in the following part concerning the Joker Objection, there may be cases of epistemic agents so severely threatening to the Minimal Principles that the exclusion from certain forms of Epistemic Contribution may appear the sole available reasonable solution. However, it is of vastly greater detriment to exclude unpredictable Contributors because of poor assessor judgement than to include at the expense of the agent in question being a considerable burden on institutional (or collective) resources. The case for social epistemic exclusion *always* rests on the faulty assumption that the agent or the population can predict the future of a particular agent's bets from the history of that agent's bets. Social epistemic exclusion can, therefore, never be *fully* justified.

To explicate AST properly, a particular understanding of the difference between social epistemic inequalities and social epistemic exclusion already accounted for during the discussion of markets' epistemic properties must be reiterated. Social inequalities are to be expected in any social system, and are not in themselves to be considered a detriment. Social inequalities account for the developments of epistemic organization, normative communities, expert structures and local transience. Moreover, following Hayek, it is advisable to understand social inequalities as population-wide signals about good and bad bets (Hayek 1982), and thus instrumental in any epistemic cooperation within a large and complex population. However, a particular type of social inequalities is of detriment to the social epistemic system – the *exclusionary* social inequalities. These are the inequalities which remove agents from the social epistemic systems, reducing Epistemic Contribution, "opportunity sets – fields of action – for individual choice" (Bromley 2008, 2), and therefore both the number and the diversity of bets, which

²⁰ Aside from the a priori argument, see for instance Heinrich 2009 for the discussion on how larger populations generate more lucky errors and subsequently increase their rate of innovation.

renders them *epistemically distortive* social inequalities. The question, however, is – how can we retain social inequalities instrumental in the development of a social epistemic system and avoid exclusionary social inequalities?

The rudimentary principle of an advanced institutional design would be that the history of bets should be used to promote agents into better positions in the knowledge-producing structure, but should not be used to exclude agents from it altogether - the agents should not be socially epistemically excluded *upon bad betting*.

Thus, an account of Epistemically Non-distortive Social Epistemic Inequalities (ENSEI) should at least allow for the agent not to be deprived of material and political conditions for Epistemic Contribution upon making bad epistemic bets. As described in Chapter 2 when discussing Hayek, system-level protection of bad betters, and more broadly the understanding of bad bets as epistemic resource within the division of epistemic labour, is fundamental to institutional epistemology. Chapter 3 further discusses what constitutes social epistemic exclusion in *applied* terms - the Minimal Principles, and thus ENSEI, being instantiated *at least* in the universal provision of sustenance, epistemic resources and possibility of Epistemic Contribution.

ENSEI would then fully allow that the history of agent's bets is surely informative for the population – it is an epistemic resource, and it allows us to discover and organize reliable knowledge-producing structures. Every agent with a history of what seem to all concerned to be good epistemic bets (expert) should with each good epistemic bet gain access to the better position in the epistemic betting structure (or, the social epistemic system). The agent with history of what appears to all concerned to be bad bets (fool) should with each bad bet be denied access to a next position in the epistemic betting structure, but is also allowed to bet *at the present position* again. The "bad" agent is not excluded from the social epistemic system by *retaining institutional inclusion even* (in the unlikely case) *when devoid of any normative community membership*, and thus her position is retained. Institutional inclusion entails *at least* the provision of sustenance, epistemic resource and possibility of Epistemic Contribution, which could then be recognized as “minimal epistemic

infrastructure”. The discovery of reliable knowledge-producing organization is, however, made possible through rewarding the “good” agents with access to further betting areas, presumably *within* normative communities. This way both requirements, that social inequalities serve the social epistemic system and that epistemically distortive exclusions are avoided, are minimally satisfied.

epistemic area 1	●	
epistemic area 0	●	● ●
epistemic area -1		
time/bet	1/2 2/2	1/2 2/2
agent type	expert	fool

fig.1: ENSEI;

with epistemic area 1 comprising of agent reliance on minimal epistemic infrastructure and membership in a normative community, epistemic area 0 comprising of agent reliance on minimal epistemic infrastructure, and epistemic area -1 of agent subject to social epistemic exclusion

Finally, another strong criteria is satisfied through this arrangement. The discovery of reliable structures is conditioned on those structures forming within the “environment” of redundant disagreement. The reliable structure, the epistemically best normative strategy, can be discovered only under conditions of contest. The protection of redundant and relatively suboptimal investigators at the global level is a thus a prerequisite for the discovery of local transience in the form of the reliable normative strategies. This principal insight of institutional epistemology must be properly reflected in the governance of our search for knowledge.

2. 3. 2. 4. Conclusion

Universal social epistemic inclusion is the function of epistemically instrumental redundant normative pluralism, constitutional ignorance of any designer or assessor of the social epistemic system and agent scarcity. Community-level exclusions may be a social epistemic virtue and are necessary for the epistemic development of the large and complex population. However, if the system fails to restrain from institutional exclusion, the system and the population are less likely to attain knowledge.

2. 3. 3. Joker Objection to Minimal Principles: Particular agents may reduce the quality of the epistemic output

Certain additional clarifications and qualifications of the present understanding of redundant pluralism and universal inclusion must be provided in anticipation of the following objection to the account. It could be proposed that a population might feature normative communities which *threaten* the social epistemic system through persistent unreasonable contest. Let us call these specific epistemically suboptimal agents and these communities – Jokers.

The answer to the objection will have to be divided into three particular propositions: 1) particular Jokers may fall under the category of the normative community *with the central aim of reducing globally sustained pluralism*, and they, if recognized as such with sufficient evidence and justification, may be regarded as best excluded from certain areas of Epistemic Contribution (if not from epistemic cooperation wholesale); 2) particular Jokers are merely Contributors; 3) given OVUT and the unlikelihood that any agent will willingly remain a Joker indefinitely due to the expected tendency among agents to conform to certain normative strategies in order to reduce the uncertainty of their environment, postponing the exclusion is the best bet. The

unlikely agent which is thrown out of all normative communities in the population should not be barred from trying to enter them or forming their own communities.

2. 3. 3. 1. The Enemies of Knowledge

Firstly, particular Jokers may fall under the category of the normative community with the aim of reducing globally sustained pluralism, which, if recognized as such with sufficient evidence and conviction, may be regarded as best excluded from certain areas of Epistemic Contribution (if not from epistemic cooperation wholesale).

If the normative community enacting the persistent unreasonable contest aims at reducing the globally sustained pluralism, there may be a strong case for denying this community access to certain areas of Epistemic Contribution. The pluralist project, and the present investigation of the design and assessment of social epistemic systems, is centrally a statement against these normative communities. Their exclusion is the condition of pluralism and universal inclusion. *“Universal” inclusion is the inclusion of all except those that deny inclusion to all and provide it only to themselves.* This qualification of “universality” is reasonable and coherent - if we wish to have a set of all at t_1+t_2 , no set of all at t_1 can include those who wish to exclude some, many or all but themselves at t_2 . The totalitarians are the first enemies of knowledge.

However, the recognition of the community as engaged in such activities and the making of the case cannot be soundly done prior to the particular case, and it is the epistemic task of the population to do it. The accusation must, of course, withstand thorough contest. It is advisable to postpone the exclusion of the normative community for the possibility of it eventually revising its normative strategy, and *any rash exclusion is fundamentally more epistemically distortive.*

2. 3. 3. 2. Relevant Contributors

Secondly, particular Jokers are merely Contributors.

The unreasonable dissent is however emphatically *not* necessarily epistemically undesirable – the pluralism should be *redundant*. Unreasonable dissent elevates the likelihood of revisiting the epistemic state being dissented to, and therefore the likelihood of it being revised if suboptimal. Joker, moreover, may be highly epistemically valuable even when the Epistemic Contribution, the content of dissent, may be evidently conceptually unsound but the target of the dissent, the particular epistemic state being objected to, may be appropriately recognized as suboptimal. If the contest is evidently unreasonable to all other normative communities, this may serve the epistemic purpose of exposing the undesirable norms and forms of contest, and perhaps of exclusion from particular normative communities. However, “(s)ometimes the threat to social stability is desirable” (Sunstein 2009, 153).

But, crucially, dissent is likely to seem unreasonable, and any population of suboptimal epistemic agents will produce inferior epistemic bets. *Redundantly* normatively pluralist *must* produce them, and any viable institutional epistemology should recognize fools and jokers as serving numerous valuable epistemic functions. Finally, as Mayo-Wilson et al (2011) show, there may even be a set of group superior epistemic states which require no individually epistemically virtuous agents. The social epistemic system lacking fools and jokers may be epistemically inferior to the one in which they abound.

2. 3. 3. 3. Unpredictable Uncertainty-reducers

Thirdly, given OVUT and the unlikeliness that any agent will willingly remain a Joker indefinitely due to the expected tendency among agents to conform to certain normative strategies in order to reduce the uncertainty of their environment, postponing the exclusion is the best bet.

The Joker Objection can be said to rest on a methodological flaw of presupposing the agent assessing the population as epistemically superior to the agents of population, specifically by presupposing that the assessor has access to knowledge of the future of an agent's epistemic output, and therefore "breaking" the baseline condition of being unable to know the future with certainty and thus the Output Value Unpredictability Thesis. Even if we can be certain (which we cannot) that Joker's epistemic output is detrimental to the overall epistemic output today, a Joker today may become a productive cooperative agent, a valuable Contributor, tomorrow.

While normative communities certainly may retain exclusionary discretions, the institutional order should strongly and decisively strive to retain agents indefinitely, even *and especially* those that show dissent. The fools, jokers, stubborn losers, deviants and strangers, and our bad bets, are the stuff of redundant normative pluralism. *If and only if a particular Joker threatens their inclusion to the point of it depending on the Joker's exclusion, the Minimal Principles must hold and the particular Joker must be excluded (at least from certain possibilities of Epistemic Contribution).*

The Joker Objection however needs to be relaxed considerably. The history of bad bets cannot be used to predict the necessary future of bad bets, and the best bet of the social epistemic system is to reduce exclusion on the basis of history of bad bets. This reduction, however, takes place within a normative environment and epistemic agents are normative "creatures". With this in mind, if a Joker is understood as somebody with a history of persistent bad bets destabilizing the normative strategies of all communities in the population, thus presuming fully justified Joker status (*entailing a positing of future of epistemically detrimental bad bets likely to a degree that exclusion becomes less epistemically detrimental than inclusion*) becomes significantly more unlikely because a Joker at t1 is more likely to become a Contributor at t2. This is so because of both a) the expected evident social pressure to do so (in the form of continuous punishment by other normative communities), and, more importantly, b) expected *normative tendency* of any agent to move towards *less uncertainty*, and thus, towards complying with a certain normative strategy and subsequently towards epistemic output of certain "non-Joker" standard.

Staving the exclusion off until it becomes undeniably impossible to coexist with the Joker remains the best bet all things considered. While the agent identified as Joker may become a significant burden on the resources of an individual Contributor, the institutional resources should be able to sustain *provisionary identified* Jokers for the open possibility of them becoming Contributors - and an *institutional* arrangement is needed to deal with reducing exclusion upon *Joker-level* bad betting.

Even if the Joker is a particular normative community, and thus a particular *norm which is at complete odds with everybody else's norms*, its exclusion would be justified if this norm, and the set of agents it governs, *were to deny other agents their fundamental substantial rights to exercise their norms in accordance with the Minimal Principles*. Otherwise, a bunch of radical, inordinate and epistemically lacking agents should still be considered more beneficial to the system as mere redundant investigators exploring the most neglected path of the landscape, then would be their social epistemic exclusion, namely deprivation of sustenance and epistemic resources as well as the possibility of Epistemic Contribution.

In sum, Joker Objection rests on an unwarranted degree of certainty in judgement of a certain agent as being incapable of contributing in the future or being actively detrimental to epistemic cooperation. That said, firstly the normative community must certainly reserve a discretion to exclude (or otherwise, a normative strategy cannot be fully enacted by the community), and secondly, there must be a certain discretion of the population of recognizing certain communities and agents as having to be excluded from particular areas of Epistemic Contribution due to the magnitude and severity of their detriment to the pluralist project and thus the epistemic infrastructure. The second case should be postponed and avoided as much as possible, preferably indefinitely. It is vastly more likely that the assessor is wrong than that the Joker "deserves" even limited institutional exclusion.

2. 4. Conclusion

The comparative standard for the assessment and design of social epistemic systems is the likelihood of revising a suboptimal epistemic state. There are baseline conditions which constrain the design of a social epistemic system, and these can be regulatively posited to be – the social epistemic systems occur necessarily in populations with finite but unknown number of epistemically suboptimal normative agents. From these two methodological givens, redundant normative pluralism and universal inclusion of epistemic agents follow as Minimal Principles of design and assessment of social epistemic systems. They are argued for via Hedge Thesis (There is more chance of revision of a suboptimal epistemic state if we all have different epistemically suboptimal normative strategies than if we all have the same) for redundant normative pluralism, and Trivial Clause (There will be more of us having different suboptimal normative strategies if there is more of us and we may have different suboptimal normative strategies), Output Value Unpredictability Thesis (It cannot be known with certainty in advance which agent will contribute the revision of the suboptimal epistemic state because no epistemic agent can predict with certainty the output value of any epistemic agent) and Agents Scarcity Thesis (Given that there can never be a sufficient number of agents for the revision of the suboptimal epistemic state to be guaranteed, it is in the best interest of all agents that each agent is in the best position to produce knowledge) for universal inclusion.

Knowledge is possible only within an interactive redundant pluralism which is globally sustained and locally transient. Universal inclusion must be institutional, while normative communities may retain discretion to exclude. Those agents and communities which threaten global pluralism may be severely punished, and those agents which pursue deeply bad strategies may be excluded by communities - thus, for example, Nazis can never come to power, and anti-vaxxers may be shunned and their norms rejected by the medical and political community as well as the community of parents and simply the community of people who care about their health. To both foster the protection of redundant investigators and not be distortive of the epistemic developments in the population, institutional universal inclusion should provide the minimal epistemic infrastructure of access to sustenance,

epistemic resources and possibility of Epistemic Contribution to all agents. Minimal Principles do not imply the conservation of all norms and an endless tie between all possible strategies. Minimal Principles must be understood as the institutional protection of conditions that allow for disagreement and the development of alternative strategies. In the following chapter I will argue that they must be enforced minimally as an infrastructure providing access to sustenance, epistemic resources and possibility of Epistemic Contribution.

The population which is redundantly normatively pluralist and universally inclusive (following HT, TC, OVUT and AST), given MECS and RBEC, is necessarily epistemically superior to any other. If this is the case, it may be argued that the present global institutional arrangements, mostly due to their exclusionary policies, developments or accidents, do not yet serve populations properly with regards to making them most likely to attain knowledge. As problems pile up, and as their complexity, interconnectedness and severity increase, the relevance of the claims advanced here should become apparent.

3. PRELIMINARY NOTES ON THE APPLIED INSTITUTIONAL EPISTEMOLOGY: EPISTEMIC INFRASTRUCTURE INDEX AND SUPPLEMENTARY ASSESSMENT

3. 1. Introduction

This final chapter seeks to open a discussion on developing an *applied* institutional epistemology which could build on the knowledge contained both in the foundational argument and the larger history of the discipline. It appears to be in the interest of the population and epistemic agents to have information on how their institutional order is apt to *epistemically* perform. Even if the version of the framework for such an assessment proposed here is insufficient, the very question of such an assessment should be regarded as a relevant enterprise. While there exist a variety of possible projects in applied institutional epistemology, I will here focus on a particular one - the assessment of minimal epistemic infrastructure, defined as provision of universal access to sustenance, epistemic resources and possibility of Epistemic Contribution. Epistemic infrastructure tracks redundant normative pluralism as the function of universal inclusion.

The structure of this chapter is as follows. I will first provide a rationale for an index which would track the minimal epistemic infrastructure. The notion was already introduced in the previous chapters, but since I believe it could provide considerable policy-making inputs and thus represent a valuable contribution by applied institutional epistemology to solving real-world problems, a reconstruction of its rationale is appropriate. I will then outline the indicators and offer further comments on them. Lastly, I will address the possible subsequent steps of an assessment, a) the appropriate research setting and agenda with regards to the analysis of results of the Index and b) the further investigation into epistemic practices and conditions through testing on agent-based models the hypotheses on epistemic distortions recognized in the population under observation.

3. 2. Rationale for the Epistemic Infrastructure Index

I will focus here on a particular problem in institutional epistemology, and discuss a policy tool for assessing the state of an social epistemic system with regards to solving that problem in applied institutional epistemology (AIE). Social epistemic system is, to repeat, an *institutional* arrangement governing over a large and normatively complex population of epistemic agents - and thus, the institutional arrangement is being assessed. The problem I will focus on is institutional protection of agents against social epistemic exclusion due to their bad betting.

Social epistemic exclusion due to arbitrary social luck is indefensible in institutional epistemology because it would basically entail exclusion of an epistemic agent from the social epistemic system without any epistemically sound reason. Bad betting may appear to be a sound reason for an exclusion of an agent - but if we regard the redundant normative pluralism epistemically instrumental, particularly if we believe it is so because epistemic agents are necessarily suboptimal, it is evident bad betting by itself cannot stand as a particularly sound epistemic reason for exclusion. This appears to be a bit of a conundrum primarily because certain epistemic inequality is desirable in a social epistemic system. The normative communities must retain exclusionary discretion to a large degree, or otherwise they could not be regarded as normative at all, and moreover, a population of such communities could hardly be called pluralist. The universal inclusion in normative communities is an incoherent idea reflective of a flat epistemology, in which any possible Epistemic Contribution is as valuable as any other. The level of inclusion relevant for the maintenance of redundant pluralism in the populations are not communities, but institutions - it is the *system-level* inclusion that should be tracked, and which I refer to as social epistemic inclusion.

Social epistemic inclusion could be understood as minimally access to sustenance (or, freedom from poverty), epistemic resources and possibility of Epistemic Contribution (primarily as access to market, voting right, free speech, free use of epistemic materials, and freedom from institutional epistemic injustice). These minimal conditions of material and political chance of a social epistemic uptake of the exhibited capability of Epistemic Contribution must be satisfied to consider an

epistemic agent included at the level of the system.

While epistemic resources and a possibility of Epistemic Contribution may appear more apparently required than freedom from poverty, this is a particularly elementary epistemic condition for any plausible epistemic progress. Poverty is epistemically detrimental inasmuch as it violates Agent Scarcity Thesis, and removes the agent from the best position to produce and find knowledge. Moreover, real poverty is largely a multidimensional and generates “corrosive” disadvantage (Alkire and Foster 2011; Wolff and de-Shalit 2013) - namely, it is an interconnected bundle of biophysical, cognitive, social, political and epistemic deprivations, which is to a relevant degree a result of epistemically detrimental social and institutional arrangements, and not a material destiny. Poverty is epistemically unjustified. Given that agents are scarce, not protecting agents from poverty is a deeply suboptimal bet at solving a highly epistemically relevant institutional problem - utilizing the dispersed knowledge in the population. It is so *even* when the person, the epistemic agent in question, makes bad bets. Withholding food, water, shelter, energy or healthcare because somebody made a bad bet is not an epistemically sound punishment. It clearly violates redundant normative pluralism and Agent Scarcity Thesis. It moreover disincentivizes further betting, and not only that specific bet.

Communities then may exclude, and some are epistemically reliable in their role in the division of labour. Certain communities may continue making bad bets, and certain stop. Some normative communities may have a better reason according to certain context- and domain-specific standard of the use of knowledge - and their reasons and epistemic procedures are an area of discussions in social and institutional epistemology. The details of finer interplays of the behaviour and structure of observed and possible normative communities, especially various political, scientific, media, health and social communities, and their relationships, the procedures of their disagreement, decision-making, problem-solving, coordination, communicating unique information, diversity and investigation, are a vast area of research in social and institutional epistemology, and their applied versions. However, no matter in which way does the normative (and institutional) setups need fixing, universal epistemic inclusion is the minimal, rudimentary *institutional* problem that must be solved by a population in order to increase the likelihood of revision of a

suboptimal epistemic lock-in.

I will now introduce and offer comments on the preliminary indicators of universal epistemic inclusion.

3. 3. Epistemic Infrastructure Index

If institutional arrangements were competing for epistemic superiority, they would certainly be advised to track the ability of social epistemic systems to provide access to sustenance, epistemic resource and possibility of Epistemic Contribution. The preliminary sketch of the indicators tracking these target conditions is as follows.

The universal access to sustenance should track minimally food, water, energy, shelter and healthcare. The universal access to epistemic resources and possibility of Epistemic Contribution should track minimally, basic liberties: freedom of speech, media, association, assembly and trade, voting rights, political freedoms, freedoms from arbitrary imprisonment; free access to educational and informational materials: free school, free university, free internet, public library, strong open access and open source policies, and limited intellectual property; certain conditions of low-cost market entry:

lack of formal denial of market entry on arbitrary grounds, free access to limited intellectual property rights for low-status contributors, antitrust laws and laws against abuse of dominance; and access to discretionary time in order for the agents to be able to diversify their epistemic contribution. These, shown in fig. 2, could be regarded as a *germinal* Epistemic Infrastructure Index (EII).

Epistemic Infrastructure Index should assess provision of universal inclusion of epistemic agents in order to indirectly assess epistemic *hedging and learning capacities* of the population (epistemically valuable diversification), and thus the likelihood of revision of suboptimal epistemic state. It tracks Minimal Principles by tracking pluralism as the *function* of resisting social epistemic exclusion of epistemic

agents upon bad betting which is, in turn, tracked through the aforementioned indicators of material and political conditions for Epistemic Contribution.

table 1: minimal Epistemic Infrastructure Index

target condition	indicators
the universal access to sustenance	food, water, energy, shelter and healthcare
the universal access to epistemic resources and possibility of Epistemic Contribution	<p>basic liberties</p> <ul style="list-style-type: none"> ● freedom of speech ● freedom of association and assembly freedom of media ● freedom of trade ● voting rights ● political freedoms ● freedom from arbitrary imprisonment <p>free access to educational and informational materials</p> <ul style="list-style-type: none"> ● free school ● free university ● free internet ● public library ● open access and open source policies ● limited intellectual property <p>low-cost market entry</p> <ul style="list-style-type: none"> ● lack of formal denial of market entry on arbitrary grounds ● free access to limited intellectual property rights for low-status

	<p style="text-align: center;">contributors</p> <ul style="list-style-type: none"> ● antitrust laws ● laws against abuse of dominance <p>access to discretionary time</p>
--	---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

The proposed set of indicators is certainly open to expansion and detailed development. Any such effort should reflect the insights of IE when understanding epistemic “infrastructure” as a network of institutional conditions conducive to revision of collective suboptimal epistemic state.

The term “infrastructure”, when denoting the group of indicators as described, could require additional clarification. All manner of devices for transgenerational transit and particularly *storing* of knowledge could be regarded as “epistemic infrastructure”, and have been, for instance by Margaret Hedstrom and John Leslie King (2006). But libraries and internet are only a part of the epistemic infrastructure. In the case of a social epistemic system as I have been discussing it, and as institutional epistemology may clearly recognize it, search for knowledge, over and above storing and transmitting, requires *extracting* it from agents. The social epistemic system needs *moves* from agents within the epistemic environment which allow for better local peaks, and thus, also, devices for storing and transmitting them. Infrastructure supporting this system should allow for *the greatest harvest of epistemic input*. As Fricker recognizes, it is the conditions for the capability to epistemically contribute, to *give*, which needs to be nurtured by an institutional arrangement. Libraries and internet are tools, necessary but insufficient to regard the agent, as Agent Scarcity Thesis would have it, *at the best position to produce knowledge*. To consider agents epistemically included in a social epistemic system, the epistemic infrastructure must account for sufficient conditions for them *not to experience severe punishment* for making an Epistemic Contribution considered relatively suboptimal, and redundant, by the most powerful coalition of other agents within a population. In order to properly understand epistemic inclusion, one needs to consider a case of transgenerational poverty, the corrosive and multidimensional disadvantage (Alkire

and Foster 2011; Wolff and de-Shalit 2013), due to both social and epistemic (and thus, crucially, *reliably correct*) prejudice. If agents do not have access to sustenance, epistemic resources and possibility for Epistemic Contribution because they are considered weird, wrong, stupid, crazy or eccentric - let alone “*unwashed*” - this is a particularly *epistemically detrimental institutional arrangement*. If Minimal Principles are pluralism and universal inclusion, both epistemically instrumental because they make the population less likely to get stuck at a suboptimal normative strategy, the elementary epistemic infrastructure should focus on providing them.

The elementary form of exclusion from the betting pool, the search for knowledge, among *human* epistemic agents is lack of access to food, water, shelter, energy and healthcare - the set of goods and services which I will refer to as *sustenance*. To be indefinitely in lack of sustenance is a punishment (Kendrex 2015). While multidimensional poverty may be used for assessing poverty levels, and further work on the precise formulation of sustenance is certainly desirable, the minimal version, tracking access to food, water, shelter, energy and healthcare, will suffice for primarily illustrative purposes. Additionally, the desired quality common to forms of sustenance which may be included in their assessment is their long-term viability. They are to be expected and agents must be able to make long-term plans and continue to bet with these goods guaranteed to hold. Trust in certain infrastructural stability of the provision of sustenance appears fundamental to its effect (Mani et al 2013) - if you are allowed sustenance occasionally, infrequently, and cannot count on it, you are *deprived* of sustenance. *The expectation of sustenance sustainability* might also be surveyed in the population.

Epistemic resources are stocks of knowledge to be used. *Free school, free university, free internet, free library and open access and open source policies* may be the most modest indicators of the availability of epistemic resources to agents. Policies of *limited intellectual property* play a strong role in availability of epistemic resources for use.

As noted, in discussing *Epistemic Contribution*, Fricker makes a key distinction between receiving and giving in an epistemic game in order to clarify why making a contribution, as a social act of offering something to other agents, should be the

capability under consideration. Engaging in the exchange of reasons, being able to test beliefs and norms, is, moreover, constitutional of epistemic agency. The population of agents must have *basic liberties* to exercise any reason-exchange worth mentioning. They must have *freedom of speech, media, association, assembly, trade and voting rights*. They must be *free to act politically* and be *free from arbitrary imprisonment*. *Open source and access policies*, again, allow for use of epistemic resources. *Certain aspects of low-cost market entry* should, beside access to sustenance, be observed through *lack of formal denial of market entry on arbitrary grounds* (which would entail any legislature that denies agents the chance to get a job, have a career, seek funding, start a business, put out a product), *free access to limited intellectual property rights for low-status contributors* (so their products are not immediately hijacked by high-status contributions), *antitrust laws* and *laws against abuse of dominance*. Lastly, *access to discretionary time* makes available opportunities to use free time in any possible way, including developing “tools” for epistemic contribution in an area alternative to the currently professional (Goodin et al 2008). Moreover, it is pertinent to enjoyment of epistemic resources, engaging in various forms of collective decision-making and problem-solving, and participating in various experimental activities. If the agent has no time to engage in alternative activity to the professional one, a considerable plausible epistemic input may be withheld. In certain ways, being denied the possibility for Epistemic Contribution encapsulates both being denied sustenance and epistemic resource - however, the distinction here is practical as it allows a finer and clearer understanding and assessment of fundamental levels of social epistemic exclusion. Access to epistemic resources and sustenance also allows for a non-perfectionist (or at least *low* perfectionist) improvement of the agents’ epistemic powers.

Freedom from poverty and oppression are quintessential to collective epistemic progress. Free speech and the possibility of Epistemic Contribution are epistemically instrumental (Fricker 2015), and the infrastructure for their material and political conditions *minimally* includes these indicators. A fuller understanding of freedom from oppression might require further research into indices of *freedom from institutional epistemic injustice* (Fricker 2007; Anderson 2012; Dotson 2014), recognition of systemic determinants of “an unwarranted infringement on the epistemic agency of knowers.” (Dotson 2012, 115). If some of noted germinal

indicators are not satisfied, however, the epistemic game is to be regarded as distorted. If we are hungry and cold, cannot get information or knowledge, cannot exercise free speech and voting rights, are incarcerated without cause, are denied market entry formally or substantially by the market being distorted by illegitimate centralizations, or have no free time, our population is inadequately *positioned to attempt* at epistemic progress. Without these conditions in place, the betting pool can be expected to grow shallow, and the revision of the suboptimal state unlikely.

As noted, it is possible certain other indicators should be added to this list. These, however, are truly minimal, and germinal. If the populations were competing in the epistemic game, they would not hesitate for a second at attempting to devise ways to meet them.

The provision of sustenance and epistemic resources may be satisfied through various institutional arrangements, as long as it exists and can be expected to be sustainable. However, if the state/market dichotomy, with its limited analytical powers, is in use, it should be noted neither sustenance nor epistemic resources need to be provided by the state as some centralized body of decision-making. It does however need to exhibit a highly predictable (and accountable) behaviour of an institution. *While providers may compete, the institution must guarantee.* The Index should merely measure *are these conditions available* - and not be concerned with how are they made available. The latter is a matter for some other investigation.

EII is a very low bar of social progress of any kind, and epistemic as well. Index should however be interested primarily in low bars - it should reflect minimal standards, the rudimentary ingredients of an epistemic setting favourable to search for knowledge. Populations with EII satisfied would differ in their epistemic performance and likely be variously suboptimal. They, however, would have a *chance at development and progress* - and this is all Index should be concerned with. The Index should assess, according to the comparative standard of revision of suboptimal epistemic state, the *minimal* standard of epistemic progress, understood as *moving away* from suboptimality. The continued institutional exercise of restraint from exclusion from the large betting pool upon bad betting appears to be the best bet as to this minimal standard. It institutionalizes diversification as the function of

inclusion of a *bad* bet. There is no pluralism if only the good bets survive. EII should track the institutional guarantee of the system-level retainment of bad betters.

The institutional inclusion is in the interest of all normative communities and epistemic agents. We should retain in our system those we don't want in our communities. *Losing unique inferential apparatuses and their normative communities because they err in judgement is a policy which rests on a flawed understanding of those apparatuses.* Centralization may occur on the level of normative communities, where pluralism may be transient. Global, and large-scale, pluralism should however remain in persistence, and the maintenance of this system-level universal inclusion should be the focus of EII. The condition that the winning epistemic bet does not determine inclusion in future epistemic betting constitutionally *allows* for the possibility of a plurality of normative enterprises. While further local peaks of epistemic development should be found and thus epistemic infrastructure might have to be significantly upgraded, its elementary purpose in any population which desires epistemic success should be to provide institutional epistemic inclusion to all, and *crucially* those individuals that are excluded from all normative communities because they are weird, wrong, stupid, crazy or eccentric, let alone "unwashed".

3. 4. Supplementary Assessment: Index Results Analysis and Contingent Distortions Review

Aside from EII, the fuller and larger assessment of the epistemic infrastructure should include an analysis of the found state, or Index Result Analysis (IRA), and the research into hypotheses on particularly problematic epistemic practices recognized in the population, Contingent Distortions Review (CDR).

Index Results Analysis (IRA) should consist of an interdisciplinary analysis of system-level failures to provide social epistemic inclusion of epistemic agents as defined by EII. It is fundamentally concerned with the flaws in the development of

epistemic infrastructure and thus epistemic *hedging* capacities of the social epistemic system. The interdisciplinary team, consisting possibly of IAD, organizational science, philosophy of society, political economy, cognitive anthropology, media studies, machine learning, and various other disciplines, should provide a diversity of perspectives on systemic relations in a task of identifying *clear points of failure to enact provision as well as the elaboration of reasons and conditions of these failures*. Their overall assessment should therefore include the account of social and institutional circumstances which hinder the development of epistemic infrastructure in the case in point. They should not assess how are people going to be smarter but why don't they have food.

Contingent Distortions Review (CDR) should, on the other hand, comprise of testing the directly epistemological diagnosis (judgements of epistemically distortive social and institutional arrangements in the case under study) by an interdisciplinary team on of agent-based models and simulations, in order to expand the understanding of plausibly particularly epistemically problematic social and institutional circumstances. Possible distortions may include: lack of interactivity and cooperation, persistence of systematic error in judgement and lack of relevant disagreement on the matter, particular configurations of normative communities, disincentivization of sharing of unique information, particular non-epistemic social conditions preventing epistemic development, and numerous other hypotheses derived from the bodies of knowledge available to the interdisciplinary team performing the Review. This part of the assessment aims at testing of sound epistemological claims on the real-world epistemic circumstances for the advancement of understanding of epistemic situations. It should expand the knowledge in applied institutional epistemology.

3. 5. Conclusion

In sum, EII focuses on an exclusive feature of the social epistemic system as a *function* of Minimal Principles - the resistance of the system to social epistemic exclusion of epistemic agents. IRA, the analysis of its results aims to comment on

the failures to enact this resistance properly. The last part of the assessment, CDR, allows for the wider epistemological review of the case under study, thus enabling the research on complex social epistemic circumstances to take place as part of the assessment. While EII should be developed as to provide a set of facts on the state of epistemic infrastructure, the Analysis and Review should be conducted by interdisciplinary teams, to reflect the advised pluralist epistemic environment for assessment.

Resisting social epistemic exclusion upon bad betting is the minimal requirement for the epistemically superior social system. The understanding which social and institutional arrangements hinder this resistance has a clear operational goal, and the possibility of testing hypotheses derived from observed epistemic settings in agent-based simulations provides the state-of-the-art research opportunities plausibly informative for further development of the discipline and the fragile and troublesome enterprise of search for knowledge under the conditions of epistemic suboptimality.

CONCLUSION

What is the problem of institutional epistemology?

Epistemic agents are necessarily suboptimal. Suboptimal agents are those that a) don't have access to all relevant evidence, b) have limited and lacking conceptual resources, c) make inferential mistakes and errors (systematic mistakes), d) tend to conserve suboptimal strategies in the search for knowledge, and e) cannot predict the future. In Chapter 2 I have argued that neither trying to make the agents less suboptimal by focusing on their individual epistemic virtues nor delegating the totality of epistemic labour to those closer to truth can fully describe a comparative standard if IE considers agents, as it should, as epistemically suboptimal. Which population is more likely to revise a suboptimal normative strategy, and thus learn, is a more appropriate, regulative and modest, comparative standard. The population which learns best is most likely to discover optimal solutions to presently unknown problems.

Thus, the problem of institutional epistemology is - *under which institutional arrangement is a large and normatively complex population of suboptimal epistemic agents least likely to conserve a suboptimal strategy in the search for knowledge?* I have argued two conditions for collective epistemic progress, Minimal Principles of IE, follow from this understanding of the institutional epistemological problem: 1) redundant pluralism and 2) universal inclusion. Since we are epistemically suboptimal, we are less likely to conserve a suboptimal epistemic strategy if we follow different strategies. And we are more likely to follow different strategies if we are not punished by social exclusion for following them. Following John Stuart Mill, I have moreover argued that this understanding satisfies the criteria of knowledge being justified true belief - knowledge is conditioned on redundant disagreement because free possibility of contesting a true belief renders that belief justified. Social epistemic inequalities can be justified only under the conditions of universal epistemic inclusion - communities may exclude me only when I can be free of poverty, have access to epistemic materials and may try to pursue their epistemic goals.

Universal access to sustenance, epistemic resources and possibility of Epistemic Contribution is the function of redundant pluralism, and thus the condition of collective epistemic progress. As I have shown in Chapter 1, institutional epistemology begins with the discussions on the utility of experimentalism, decentralization, diversity and division of epistemic labour for purposes of social learning. The findings of these investigations find their true expression in minimal epistemic infrastructure. If our social epistemic systems fail to protect redundant investigators, we are less likely to attain knowledge.

References

- [1] Aligica, P. D. 2014. *Institutional Diversity and Political Economy: The Ostroms and Beyond*. New York: Oxford University Press.
- [2] Aligica, P. D. and Tarko, V. 2013. Co-Production, Polycentricity, and Value Heterogeneity: The Ostroms' Public Choice Institutionalism Revisited. *American Political Science Review* 107/4: 726-741.
- [3] Aligica, P. D., Boettke, P. J. & Tarko, V. 2019. *Public Governance and the Classical-Liberal Perspective*. New York: Oxford University Press.
- [4] Alkire S. and Foster J. 2011. Counting and multidimensional poverty measurement. *Journal of Public Economics* 95/7-8: 476-487.
- [5] Anderson, E. 2006. The Epistemology of Democracy. *Episteme: A Journal of Social Epistemology* 3/1-2: 8-22.
- [6] Anderson, E. 2012. Epistemic Justice as a Virtue of Social Institutions. *Social Epistemology: A Journal of Knowledge, Culture and Policy* 26/2: 163-173.
- [7] Axelrod, R. 1986. An Evolutionary Approach to Norms. *The American Political Science Review* 80/4: 1095-1111.
- [8] Berg, N. 2003. Normative behavioral economics. *Journal of Socio-Economics* 32: 411–427.
- [9] Bicchieri, C. 2006. *The Grammar of Society: The Nature and Dynamics of Social Norms*. New York: Cambridge University Press.
- [10] Bicchieri, C. 2017. *Norms in the Wild: How to Diagnose, Measure, and Change Social Norms*. Oxford: Oxford University Press.

- [11] Boettke, P. J., Caceres, W. Z. & Martin, A. G. 2013. Error is Obvious, Coordination is the Puzzle. In *Hayek and Behavioral Economics*, ed. R. Frantz and R. Leeson, 90-110. New York: Palgrave MacMillan.
- [12] Boettke, P. J., Tarko V. and Aligica P. 2016. Why Hayek Matters: The Epistemic Dimension of Comparative Institutional Analysis. In *Revisiting Hayek's Political Economy*, ed. P. J. Boettke and V. H. Storr, 163-185. Emerald Group Publishing Limited.
- [13] Boettke P. 2018. *F. A. Hayek: Economics, Political Economy and Social Philosophy*. London: Palgrave Macmillan.
- [14] Brandom, R. 2001. *Making It Explicit: Reasoning, Representing and Discursive Commitment*. Cambridge and London: Harvard University Press.
- [15] Brandom, R. 1998. Action, Norms and Practical Reasoning. *Philosophical Perspectives* 12: 127-139.
- [16] Bromley, D. 2008. Beyond Market Failure: Volitional Pragmatism as a New Theory of Public Policy. *Economia Politica* 25/2: 219-241.
- [17] Caldwell, B. and Reiss, J. 2006. Hayek, Logic, and the Naturalistic Fallacy. *Journal of the History of Economic Thought* 28/3: 359-370.
- [18] Case, S. 2016. Normative Pluralism Worthy of the Name is False. *Journal of Ethics and Social Philosophy*, 11/1.
- [19] Cerovac, I. 2018. Epistemic Liberalism. *Prolegomena* 17/1: 81-95.
- [20] Condorcet, M. J. A. N. 1976. *Condorcet: Selected Writings*. Indianapolis: Bobbs-Merrill.
- [21] D'Agostino, F. 2009. From Organization to Division of Cognitive Labor. *Politics, Philosophy & Economics* 8/1: 101-129.

- [22] Danon, L. 2019. Animal Normativity. *Phenomenology and Mind* 17: 176-187.
- [23] Dewey, J. 1976. Valuation and Experimental Knowledge. In *The Middle Works of John Dewey, 1899–1924, vol. 13*, ed. J. A. Boydston, 3–28. Carbondale: Southern Illinois University Press.
- [24] Dewey, R. 1977. *The Philosophy of John Dewey: A Critical Exposition of his Method, Metaphysics, and Theory of Knowledge*. Hague: Martinus Nijhoff.
- [25] Dotson, K. 2014. Conceptualizing Epistemic Oppression. *Social Epistemology: A Journal of Knowledge, Culture and Policy* 28/2: 115-138.
- [26] Esposito, A. 2014. Path Dependence, Lock-in, and Social Norms. *International Journal of Economic Issues* 7/1: 1-15.
- [27] Estlund, D. 2008. *Democratic Authority: A Philosophical Framework*. Princeton: Princeton University Press.
- [28] Estlund, D. 2003. Why Not Epistocracy? In *Desire, Identity and Existence: Essays in honor of T. M. Penner*, ed. N. Reshotko, 53-69. Academic Printing and Publishing.
- [29] Foa R. S. and Mounk Y. 2016. The Danger of Deconsolidation: The Democratic Disconnect. *Journal of Democracy* 27/3: 5-17.
- [30] Fossen, T. 2014. Politicizing Brandom's Pragmatism: Normativity and the Agonal Character of Social Practice. *European Journal of Philosophy* 22/3: 371-395.
- [31] Fleetwood, S. 1995. *Hayek's Political Economy: The socio-economics of order*. London and New York: Routledge.

- [32] Fricker, M. 2007. *Epistemic Injustice: Power and the Ethics of Knowing*. New York: Oxford University Press.
- [33] Fricker, M. 2015. Epistemic Contribution as a Central Human Capability. In *The Equal Society: Essays on Equality in Theory and Practice*, ed. G. Hull, 73-91. London: Lexington Books.
- [34] Gaus, G. 2008. Is the Public Incompetent? Compared to Whom? About What? *Critical Review* 20: 291-311
- [35] Gaus, G. 2011. *The Order of Public Reason: A Theory of Freedom and Morality in a Diverse and Bounded World*. New York: Cambridge University Press.
- [36] Gaus, G. 2011. On Seeking the Truth (Whatever That Is) through Democracy: Estlund's Case for the Qualified Epistemic Claim. *Ethics* 121: 270–300.
- [37] Gaus, G. 2013. The Evolution, Evaluation and Reform of Social Morality: A Hayekian Analysis. In *F. A. Hayek and the Modern Economy*, ed. S. Levy and D. Peart, 59-88. New York: Palgrave Macmillan.
- [38] Gaus, G. 2016. *The Tyranny of the Ideal: Justice in a Diverse Society*. New Jersey: Princeton University Press.
- [39] Gaus, G. 2018a. The Complexity of a Diverse Moral Order. *The Georgetown Journal of Law and Public Policy* 16: 645-79.
- [40] Gaus, G. 2018b. Political Philosophy as the Study of Complex Normative Systems. *Cosmos + Taxis* 5/2: 62-76.
- [41] Gilbert, M. 2014. *Joint Commitment: How We Make the Social World*. New York: Oxford University Press.
- [42] Ginsborg, H. 2011. Primitive normativity and skepticism about rules. *The Journal of Philosophy* 108/5: 227-254.

- [43] Goldman, A. 1979. What is Justified Belief? In *Justification and Knowledge*, ed. G. Pappas, 1-25. Boston: D. Reidel.
- [44] Goldman, A. and Cox J. C. 1996. Speech, Truth, and the Free Market for Ideas. *Legal Theory 2*: 1- 32.
- [45] Goldman, A. 2011. A Guide to Social Epistemology. In *Social Epistemology: Essential Readings*, ed. A. Goldman and D. Whitcomb, 11-37. New York: Oxford University Press.
- [46] Goodin, R., Rice, J. M., Parpo, A. and Eriksson, L. 2008. *Discretionary Time: A New Measure of Freedom*. Cambridge: Cambridge University Press.
- [47] Gronow, A. 2012. From Habits to Social Institutions: A Pragmatist Perspective. In *The Habits of Consumption*, ed. A. Warde and D. Southerton, 26-45. Helsinki Collegium for Advanced Studies.
- [48] Guala, F. and Hindriks, F. 2015. A Unified Social Ontology. *The Philosophical Quarterly* 65/259: 177-201.
- [49] Hastie, R. and Dawes, R. 2001. *Rational Choice in an Uncertain World: The Psychology of Judgement and Decision Making*. Thousand Oakes: Sage Publications.
- [50] Hayek, F. A. 1945. The Use of Knowledge in Society. *The American Economic Review* 35/4: 519-530.
- [51] Hayek, F. A. 1960. *The Constitution of Liberty: Definitive Edition*. Chicago: The University of Chicago Press.
- [52] Hayek, F. A. 1978a. Coping With Ignorance. *Imprimis* 7(7): 1-6.

[53] Hayek, F. A. 1978b. *New Studies in Philosophy, Politics, Economics and History of Ideas*. London: Routledge

[54] Hayek, F. A. 1982. *Law, Legislation and Liberty: A new statement of the liberal principles of justice and political economy*. London: Routledge.

[55] Hayek F. A. 2014. *The Market and Other Orders*. Chicago: The University of Chicago Press.

[56] Hedstrom, M. and King J. L. 2006. Epistemic Infrastructure in the Rise of the Knowledge Economy. In *Advancing Knowledge and the Knowledge Economy*, ed. B. Kahin and D. Foray, 113-134. Cambridge: MIT Press.

[57] Heinrich, J. 2009. Why societies vary in their rates of innovation: The evolution of innovation-enhancing institutions. In *Innovation in Cultural Systems: Contributions from Evolutionary Anthropology (Altenberg Workshops in Theoretical Biology)*, ed. S. Shennan and M. O'Brien. Altenberg: Konrad Lorenz Institute.

[58] Herrmann-Pillath, C. 2012. Institutions, Distributed Cognition and Agency: Rule-following as Performative Action. *Journal of Economic Methodology* 19/1: 21-42.

[59] Hodgson, G. 1991. Hayek's Theory of Cultural Evolution: An evaluation in light of Vanberg's critique. *Economics and Philosophy* 7: 67-82

[60] Hodgson, G. 2006. What Are Institutions? *Journal of Economic Issues* 40/1: 1-25.

[61] Hong, L. and Page. S. 2004. Groups of diverse problem solvers can outperform groups of high-ability problem solvers. *Proceedings of the National Academy of Sciences of the United States of America* 101/46: 16385-16389.

[62] Kahneman, D. 2011. *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux.

- [63] Kelly, P. 2006. Liberalism and Epistemic Diversity: Mill's Sceptical Legacy. *Episteme* 3/3: 248-263.
- [64] Kendrex, J. 2015. Punishing the Poor Through Welfare Reform: Cruel and Unusual? *Duke Law Journal* 64: 121-140.
- [65] Kingston, C. and Caballero, G. 2009. Comparing Theories of Institutional Change. *Journal of Institutional Economics* 5/2: 151-180.
- [66] Kirzner, I. 1997. Entrepreneurial Discovery and the Competitive Market Process: An Austrian Approach. *Journal of Economic Literature* 35/1: 60-85.
- [67] Kitcher, P. 1990. Division of Cognitive Labor. *The Journal of Philosophy* 87/1: 5-22.
- [68] Klein, D., Marx, J. and Fischbach, K. 2018. Agent-Based Modeling in Social Science, History, and Philosophy: An Introduction. *Historical Social Research* 43/1: 7-27.
- [69] Krstić, M. 2012. The role of rules in the evolution of the market system: Hayek's concept of evolutionary epistemology. *Economic Annals* 57/194: 123-140.
- [70] Landemore, H. 2012a. Why the Many Are Smarter than the Few and Why It Matters. *Journal of Public Deliberation* 8/1: article 7.
- [71] Landemore, H. 2012b. Democratic Reason: The Mechanism of Collective Intelligence in Politics. In *Collective Wisdom: Principles and Mechanisms*, ed. H. Landemore and J. Elster, 251-290. New York: Cambridge University Press.
- [72] Landemore, H. 2013. *Democratic Reason: Politics, Collective Intelligence, and the Rule of the Many*. Princeton: Princeton University Press.
- [73] Lehrer, K. 2000. *Theory of Knowledge*. Boulder: Westview.

[74] List, C. and Pettit, P. 2011. *Group Agency: the Possibility, Design, and Status of Corporate Agents*. Oxford: Oxford University Press.

[75] Mani, A., Mullainathan S., Shafir, E., Zhao, J. 2013. Poverty Impedes Cognitive Function. *Science* 341/6149: 976-980.

[76] Mann, A. 2007. The Power of Prediction Markets. *Nature* 538: 308-310.

[77] Mayo-Wilson, C., Zollman, K. and Danks, D. 2011. The Independence Thesis: When Individual and Social Epistemology Diverge. *Philosophy of Science* 78/4: 653-677

[78] Mercier, H. and Sperber, D. 2017. *The Enigma of Reason*. Cambridge: Harvard University Press.

[79] Mill, J. S. 2003. On Liberty. In *Utilitarianism and On Liberty*, 88-181. Oxford: Blackwell Publishing.

[80] Moyal-Sharrock, M. 2016. The Animal in Epistemology: Wittgenstein's Enactivist Solution to the Problem of Regress. *International Journal for the Study of Skepticism* 6: 97-119.

[81] Muldoon, R. and Weisberg, M. 2011. Robustness and Idealization in Models of Cognitive Labor. *Synthese* 183/2: 161–174.

[82] Muldoon, R. 2013. Diversity and the Division of Cognitive Labour. *Philosophy Compass* 8/2: 117-125.

[83] North, D. 1990. *Institutions, Institutional Change, and Economic Performance*. Cambridge: Cambridge University Press.

[84] North, D. and Denzau, A. 1994. Shared Mental Models: Ideologies and Institutions. *Kyklos* 47/1: 3-31.

[85] Novak, M. 2018. *Inequality: An Entangled Political Economy Perspective*. Cham: Palgrave Macmillan.

[86] O'Driscoll, Jr., G. P. and Rizzo, M. J. 1985. *The Economics of Time and Ignorance*. New York: Basil Blackwell Ltd.

[87] Ostrom, E. 1990. *Governing the Commons: The Evolution of Institutions for Collective Action*. New York: Cambridge University Press.

[88] Ostrom, E. 2000. Collective Action and the Evolution of Social Norms. *The Journal of Economic Perspectives* 14/3: 137-158.

[89] Ostrom, E. 2005. *Understanding Institutional Diversity*. Princeton and Oxford: Princeton University Press.

[90] Ostrom, E. and Hess, C. 2007. A Framework for Analyzing the Knowledge Commons. In *Understanding Knowledge as a Commons*, ed. E. Ostrom and C. Hess, 41-83. Cambridge: The MIT Press

[91] Ostrom, E. 2010. Beyond Markets and States: Polycentric Governance of Complex Economic Systems. *The American Economic Review* 100/3: 641-672.

[92] Page, S. 2008. *The Difference: How the Power of Diversity Creates Better Groups, Schools, Firms, and Societies*. Princeton: Princeton University Press.

[93] Peter, F. 2013. Epistemic Foundations of Political Liberalism. *Journal of Moral Philosophy* 10/5: 598-620.

[94] Peter, F. 2016. The Epistemic Circumstances of Democracy. In *The Epistemic Life of Groups: Essays in the Epistemology of Collectives*, ed. M. S. Brandy and M. Fricker, 133-149. Oxford: Oxford University Press.

[95] Pluchino, A., Rapisadra A., Garofalo C. 2010. The Peter Principle Revisited: A Computational Study. *Physica A* 389: 467-472.

- [96] Pritchard, D. 2005. *Epistemic Luck*. Oxford: Oxford University Press.
- [97] Pritchard, D. 2011. Epistemic Relativism, Epistemic Incommensurability, and Wittgensteinian Epistemology. In *A Companion to Relativism*, ed. S. Hales, 266-285. Hoboken: Wiley-Blackwell.
- [98] Reijula S. and Kuorikoski, J. 2020. Modeling epistemic communities. In *The Routledge Handbook of Social Epistemology*, ed. M. Fricker, P. J. Graham, D. Henderson, N. Pedersen, J. Wyatt, 240-250. New York: Routledge.
- [99] Rietveld, E. 2008. Situated Normativity: The Normative Aspect of Embodied Cognition in Unreflective Action. *Mind* 117/468: 973-100.
- [100] Rittel, H. and Webber, M. 1973. Dilemmas in a General Theory of Planning. *Policy Sciences* 4: 155-169.
- [101] Roth, A., Prasnikar, V., Okuno-Fujiwara, M., and Zamir, S. 1991. Bargaining and Market Behavior in Jerusalem, Ljubljana, Pittsburgh and Tokyo: An Experimental Study. *The American Economic Review* 81/5: 1068-1095.
- [102] Rothstein, B. 2011. *The Quality of Government: Corruption, Social Trust, and Inequality in International Perspective*. Chicago and London: The University of Chicago Press.
- [103] Schwartzberg, M. 2015. Epistemic Democracy and Its Challenges. *Annual Review of Political Science* 18: 187-203.
- [104] Sellars, W. 1963. Empiricism and the Philosophy of Mind. In *Science, Perception and Reality*, 127-197. Atascadero: Ridgeview Publishing Company.
- [105] Sunstein, C. 1996. Social Norms and Social Roles. *Columbia Law Review* 96/4: 903-968.

[106] Sunstein, C. 2006. *Infotopia: How Many Minds Produce Knowledge*. New York: Oxford University Press.

[107] Sunstein, C. 2009. *Going to Extremes: How Like Minds Unite and Divide*. New York: Oxford University Press.

[108] Talisse R. 2009. *Democracy and Moral Conflict*. Cambridge: Cambridge University Press.

[109] Tarko, V. 2015. The Role of Ideas in Political Economy. *Review of Austrian Economics* 28: 17–39.

[110] Thompson, A. 2014. Does Diversity Trump Ability? An Example of Misuse of Mathematics in Social Sciences. *Notices of the AMS* 61/9: 1024-1030.

[111] Tuomela, R. 2013. *Social Ontology: Collective Intentionality and Group Agents*. Oxford: Oxford University Press.

[112] Unger, P. 1968. An Analysis of Factual Knowledge. *Journal of Philosophy* 65: 157-70.

[113] Vanberg, V. 2006. Rationality, rule-following and emotions: on the economics of moral preferences. In *Papers on Economics and Evolution*, 0621. Jena: Max Planck Institute of Economics.

[114] Wallis, J.J. 2011. Institutions, organizations, impersonality, and interests: The dynamics of institutions. *Journal of Economic Behaviour and Organization* 79/1-2: 48-64.

[115] Weisberg, M. and Muldoon, R. 2009. Epistemic Landscapes and the Division of Cognitive Labour. *Philosophy of Science* 76/2: 225-252.

[116] Weisberg, M. 2010. New Approaches to the Division of Cognitive Labor. In *New Waves in Philosophy of Science*, ed. P. D. Magnus and J. Busch, 250-269. New York: Macmillan.

[117] Wolff, J. and de-Shalit A. 2013. *Disadvantage*. Oxford: Oxford University Press.

[118] Zollman, K. 2010. The Epistemic Benefit of Transient Diversity. *Erkenntnis* 72/1: 17-35.

Figures and tables

Fig. 1: Epistemically Non-distortive Social Epistemic Inequalities (ENSEI).....114

Table 1: Minimal Epistemic Infrastructure Index.....126

Marko-Luka Zubčić was born on July 21st 1986 in Rijeka. He has an MA in philosophy and english language and literature from Faculty of Humanities and Social Sciences in Rijeka. He is an associate of Center for Advanced Studies South East Europe. He works on policy analysis and development, communication strategies, and project management, in collaboration with various institutions and NGOs.

Marko-Luka Zubčić rođen je 21. 7. 1986. u Rijeci. Diplomirao je filozofiju i engleski jezik i književnost na Filozofskom fakultetu u Rijeci. Stručni suradnik je Centra za napredne studije - Jugoistočna Europa. U suradnji s nizom institucija i nevladinih organizacija radi na razvoju i analizi javnih politika, oblikovanju i provedbi komunikacijskih strategija, i upravljanju projektima.

Published papers/Objavljeni radovi:

Zubčić, ML. 2017. Belief-revision, Epistemic Contribution, and Polynormativity. *Ethics and Politics* 19/2: 207-227.

Zubčić, ML. 2019. Comparative standard in institutional epistemology. *Philosophy and Society* 30/3: 418-430, <https://doi.org/10.2298/FID1903418Z>

Zubčić, ML. 2019. Ignorance, norms and instrumental pluralism: Hayekian institutional epistemology. *Synthese*, DOI: 10.1007/s11229-019-02420-5

