



Jurisprudence

An International Journal of Legal and Political Thought

ISSN: (Print) (Online) Journal homepage: www.tandfonline.com/journals/rjpn20

A conceptual framework for legal personality and its application to AI

Claudio Novelli, Giorgio Bongiovanni & Giovanni Sartor

To cite this article: Claudio Novelli, Giorgio Bongiovanni & Giovanni Sartor (2022) A conceptual framework for legal personality and its application to AI, *Jurisprudence*, 13:2, 194-219, DOI: [10.1080/20403313.2021.2010936](https://doi.org/10.1080/20403313.2021.2010936)

To link to this article: <https://doi.org/10.1080/20403313.2021.2010936>



© 2021 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group



Published online: 09 Dec 2021.



Submit your article to this journal [↗](#)



Article views: 7857



View related articles [↗](#)



View Crossmark data [↗](#)



Citing articles: 2 View citing articles [↗](#)

A conceptual framework for legal personality and its application to AI

Claudio Novelli^{a,b,c}, Giorgio Bongiovanni^{a,b} and Giovanni Sartor^{a,b,d}

^aDepartment of Legal Studies, University of Bologna, Bologna, Italy; ^bCIRSFID – Alma AI, University of Bologna, Bologna, Italy; ^cDroit – University of Luxembourg, Luxembourg; ^dDepartment of Law, European University Institute, Florence, Italy

ABSTRACT

In this paper we provide an analysis of the concept of legal personality and discuss whether personality may be conferred on artificial intelligence systems (AIs). Legal personality will be presented as a doctrinal category that holds together bundles of rights and obligations; as a result, we first frame it as a node of inferential links between factual preconditions and legal effects. However, this inferentialist reading does not account for the ‘background reasons’ of legal personality, i.e., it does not explain why we cluster different situations under this doctrinal category and how extra-legal information is integrated into it. We argue that one way to account for this background is to adopt a neoinstitutional perspective and to update the ontology of legal concepts with a further layer, the meta-institutional one. We finally argue that meta-institutional concepts can also support us in finding an equilibrium around the legal-policy choices that are involved in including (or not including) AIs among legal persons.

KEYWORDS

Legal personhood; AI; ontology; metaphysics; concepts; grounding

1. Introduction: a conceptual framework for legal personality

The contours of the status of legal personality are of major theoretical value, considering the systemic role personality plays not only within legal system, but also within other networks of rules governing human conduct.¹ Legal personification may have important social effects, by empowering and protecting the entities personified and by working to the benefit of the general interest (e.g., by incentivising economic or social initiatives).² There is much debate about the boundaries of legal personality: about whether not only humans and their organisations, but also nonhuman animals, environmental entities, idols, unborn children, and software agents should be granted personality status.³

In most legal systems, the entities just mentioned are not seen as responsible agents. Some of them are seen as incapable of purposeful action, and others (e.g., animals), while

CONTACT Claudio Novelli  claudio.novelli@unibo.it

¹Alexander Nékam, *The Personality Conception of the Legal Entity* (Harvard University Press 1938).

²Henry Hansmann, Reinier Kraakman and Richard Squire, ‘Law and the Rise of the Firm’ (2005) 119 *Harvard Law Review* 1335.

³See, among the others, Visa Kurki and Tomasz Pietrzykowski (eds), *Legal Personhood: Animals, Artificial Intelligence and the Unborn* (Springer International Publishing 2017).

endowed with a limited agency, are seen as incapable of reasoning or deliberating and devoid of moral standing. But it has not always been like that: there have been times and places where, for instance, nonhuman animals have been put on trial and charged with proper crimes (and even convicted).⁴

Conversely, there have been contexts, e.g., Roman law, in which the status of a full-fledged person was only granted to certain human individuals, to the exclusion of others (e.g., slaves).⁵ Modern legal systems, on the other hand, include all human beings in the community of legal persons because only humans are deemed to be morally and mentally worthy of holding legal positions. Personality is also extended to various types of human organisations – operating in the economic, social, or religious sphere –, since such organisation advance human interests, to this end relying on the cognitive skills of their human representatives and agents.

The promoters of personhood for nonhuman entities, mainly in the case of animals, often claim that some of the cognitive skills typically relevant to morality and law – skills such as basic rationality, understood as the capacity to adopt goals/desires and acting toward them – are not exclusive to humans.⁶ In moral philosophy, a version of this idea is also known as the ‘the argument from marginal cases’, and it challenges the inconsistency of ascribing moral standing to marginal human beings with very limited, or no, cognitive skills – e.g., individuals who are severely mentally disabled, pre-rational (children) or post-rational (the senile) – while not doing the same for animals that meet equal or even superior rationality standards.⁷

Others have argued that the capacity for purposeful action is not an essential precondition of personality, since personification may instead be aimed at protecting entities having properties such as sentience, dignity, or vulnerability, or at advancing interests such as preserving the environment or the welfare of the community.⁸ One objection to this view might be that personality is neither the only available nor the most appropriate legal instrument with which to protect these entities or advance these interests.

These disputes rest on different understandings of the roots of legal personality, as well as of the overall legal practice around it, and reveal that the parameters used by legal systems to ascribe personality are still controversial and evolving.

To provide a context for the debate on legal personality, we will address legal personality in a comprehensive manner, by means of a conceptual map/ontology showing the relationships between the norms of positive law governing this status and the deeper set of reasons justifying them. We will indeed provide a model of the conceptual structure of legal personality – a model which may potentially be applied to other central legal concepts, and which remains as neutral as possible with respect to specific philosophical conceptions, e.g., monist or legalist (Section 3).

⁴Edward Payson Evans, *The Criminal Prosecution and Capital Punishment of Animals* (William Heinemann 1906).

⁵William J Curran, ‘An Historical Perspective on the Law of Personality and Status with Special Regard to the Human Fetus and the Rights of Women’ (1983) 61 *The Milbank Memorial Fund Quarterly. Health and Society* 58.

⁶On animal rights see, in the first place, Peter Singer, *Animal Liberation* (Harper Collins 1975); more specifically on the autonomy of animals and legal status in Steven Wise, *Rattling the Cage: Towards Legal Rights for Animals* (Perseus Publishing 2000).

⁷For an introduction to the subject see, among others, Daniel Dombrowski, *Babies and Beasts: The Argument from Marginal Cases* (The University of Illinois Press 1997). On a more specific issue Julia K Tanner, ‘The Argument from Marginal Cases and the Slippery Slope Objection’ (2009) 18 *Environmental Values* 51.

⁸A comprehensive analysis of these profiles can be found in Joshua C Gellers, *Rights for Robots Artificial Intelligence, Animal and Environmental Law* (Routledge 2021).

Some notions of analytical metaphysics – mainly *grounding* – will be used to illustrate the (constitutive) relations that link certain facts and properties to legal kinds (Section 4). It will be argued that legal personality is *grounded* both in a set of facts and properties described by legal rules and in a set of facts and properties operating in the background.

An approach of social ontology will be advocated which has a prominent counterpart in the theory of law, and that is so-called neoinstitutionalism; in this context, legal personality will be characterised as an institutional legal concept, following Neil MacCormick's vision; and in partial opposition to an inferentialist reading (Sections 6 and 7).

We shall construct for legal concepts a multilayered ontology consisting of an institutional, an intermediate, and a meta-institutional layer. Each layer describes sets of facts, values, and properties relevant for the personality status. The content of the institutional layer is fixed by legal rules, while the other two layers tend to be extra-legal, in a sense that we will later explain (Sections 8 and 9). The intermediate level includes concepts that, like midlevel terms in morality, can help to generate reflexive equilibria in doctrinal and legislative debates, facilitating theoretical agreements under circumstances of indeterminacy (Section 10).

Each nonhuman entity's eligibility (for personhood) deserves its own in-depth study, since different considerations may apply. For instance, the way an animal might benefit from personality is quite different from the way a river might.⁹

Here we focus on the idea of conferring legal personality on a peculiar type of artefact: artificial intelligence systems. We start out with an analysis of the debate on conferring legal personality on AI systems (Section 2). This analysis will enable us to discern the nature of the debate on legal personhood and the presupposition under which it takes place. It will therefore set the stage for the subsequent conceptual analysis (Section 10). Our inquiry does not have a purely theoretical purpose but is also driven by a practical need. Indeed, the 'noise' surrounding the idea of AI personality requires clarification on (a) what are the assumptions and modes of existence of legal personality, and (b) whether alternative strategies are available by which the information and rules that tend to be associated with personality can be organised in such a way as to facilitate convergence among different views.

2. AI and law

More and more often humans engaged in different activities – as users of a self-driving vehicles, consumers making contracts with bots, physicians supported by diagnostic systems – are interacting with artificial entities that play an intelligent and active role.

Many definitions of artificial intelligence exist.¹⁰ A broad enough, and often cited one, was proposed by the leading AI scientist John McCarthy, according to whom AI 'is the science and engineering of making intelligent machines, especially intelligent computer programs. It is related to the similar task of using computers to understand human

⁹A traditional reading here is Christopher D Stone, 'Should Trees Have Standing? Toward Legal Rights for Natural Objects' (1972) 45 Southern California Law Review 450. There have also been some concrete initiatives to promote the attribution of so-called 'environmental personality', as the recent New Zealand experience proves; see here Aikaterini Argyrou and Harry Hummels, 'Legal Personality and Economic Livelihood of the Whanganui River: A Call for Community Entrepreneurship' (2019) 44 Water International 752.

¹⁰Stuart Russell and Peter Norvig, *Artificial Intelligence: A Modern Approach, Global Edition* (Pearson Education 2016)

intelligence, but AI does not have to confine itself to methods that are biologically observable'.¹¹ A more structured definition has been recently proposed by the expert group appointed by the European Commission in 2019 (AI HLEG), which describes artificial intelligence systems (AIs) as

software (and possibly also hardware) systems designed by humans that, given a complex goal, act in the physical or digital dimension by perceiving their environment through data acquisition, interpreting the collected structured or unstructured data, reasoning on the knowledge, or processing the information, derived from this data and deciding the best action(s) to take to achieve the given goal. AI systems can either use symbolic rules or learn a numeric model, and they can also adapt their behaviour by analysing how the environment is affected by their previous actions [...].¹²

Such systems can be put to multiple uses, both with and without hardware or robotic components, to build technologies such as driverless cars, autonomous weapons, medical diagnosis systems, and contract agents.

It must be specified that most AI systems only perform a fraction of the activities listed in the definition by the HLEG: pattern recognition (e.g., recognising images of plants, animals, objects, human faces, or attitudes), language processing (e.g., understanding spoken languages, translating between languages, fighting spam, or answering queries), practical suggestions (e.g., recommending purchases, purveying information, performing logistic planning, or optimising industrial processes), etc. On the other hand, some systems may combine many such capacities, as in the example of self-driving vehicles or military and care robots.

Since our analysis is focused on granting legal personality to AI systems, we shall only consider artificial agents that have an ability to pursue goals autonomously, meaning that while their top-level goals may be defined externally (by the agents' designers and users), the way in which such top-goals are pursued by the agents – including the choice of lower-level instrumental goals – is responsive to the physical and social environment in which the agents act. Thus, for instance, even if automated translation or a machine-learning system designed to identify objects in images qualifies as an AI system, neither would fall within the scope of our analysis. Advanced digital bots used in online commerce, or physical robots providing services such as transportation or care can come within our scope as long as they process information and make choices among epistemic and practical alternatives pertaining to their tasks, make consequential changes in the digital or physical world, and engage in communication. The action of such AI systems appears to be autonomous, not being pre-programmed by human designers as a specific response to the particular circumstances at hand. It is up to such systems to process the available information, using learning and inferential algorithms, and to make assessments and decisions accordingly.

Autonomous artificial agents are increasingly able to cope with uncertain and dynamic environments, adapting to lack of information, acquiring new knowledge, and making appropriate choices. An additional complexity is owed to the fact that

¹¹ John MacCarty, 'What is Artificial Intelligence?' <<https://jmc.stanford.edu/articles/whatisai.html>> accessed 12 November 2007.

¹² Ethics Guidelines for Trustworthy AI, High-Level Expert Group on Artificial Intelligence, set up by the European Commission, 2019, p. 36.

each such system may be part of a dense network of other artificial agents, sharing information and making coordinated decisions (e.g., a fleet of autonomous cars).

Artificial agents can also participate in legally relevant interactions, producing effects in the sphere of their owners as well as that of human third parties. Their use can produce legal effects, i.e., contractual rights and obligations, for their users (e.g., vendors in electronic commerce) by virtue of their ability to interact with other subjects (human and electronic) entering into legal transactions.

This raises a set of issues pertaining to accountability and moral responsibility, as well as to legal liabilities. Who is to be called on explain and justify the behaviour of such a system? Who may be subject to blame or adverse consequences if such behaviour turns out to be harmful? Who is to be subject to legal penalties, including damages as well as administrative fines or criminal punishment, when such behaviour violates applicable legal rules? Is it the owner, the user, the programmer, the manager of the development team, the person tasked with testing and debugging? What happens if an action cannot be traced back to any particular human?

The cognitive and operational abilities acquired by AIs invite us to question whether they are properly classified under the law as mere ‘things’, rather than as legal persons, and whether this will continue to be their legal classification in the future. In fact, non-human animals have sometimes been argued to have rights on the basis that they possess human-like attributes of reasoning and ‘practical autonomy’ such as the ability to recognise cause and effect, form desires and engage in self-cognition. Some AI systems seem to possess these attributes as well, and possibly to a higher degree.¹³

The issue of whether some AI systems should acquire legal personality has been at the centre of a lively debate. Most jurists seem to prefer to address the legal issues emerging from the agency of AI systems by adapting existing liability models – i.e., to make users and producers liable for the autonomous and unpredictable actions of AI systems – rather than to extend legal personality. Although a cautious approach seems reasonable for first-generation artificial agents, this may prove to be insufficient as the autonomy of AI system increases.¹⁴ Against this background, the attempt by the European Parliament to include the ‘electronic personhood’ of robots in the political debate has not been taken up by the European Commission.¹⁵ This proposal struggles to gain traction, since some of the risks associated with the advancement of AI can be tackled with the liability rules already in force within our legal systems, or by preventive measures and controls (as in the recent EU proposal for an AI regulation). Moreover, the European Union has no competence to decide what counts as a ‘person’ within national legal systems, as such a decision rests with the Member States.¹⁶ However, we think that the time has come to consider whether a comprehensive legal approach to the autonomous behaviour of

¹³Appeal to the notion of practical autonomy is emphasised, in particular, by Steven Wise in Steven Wise, *Drawing the Line: Science and the Case for Animal Rights* (Basic Books 2003).

¹⁴Emad Abdel Rahim Dahiyat, ‘Law and Software Agents: Are They “Agents” by the Way?’ (2020) 29 *Artificial Intelligence and Law* 59.

¹⁵See the European Parliament’s resolution of 16 February 2017, with recommendations to the Commission on Civil Law Rules on Robotics and the European Commission’s subsequent 25 April 2018 outline ‘Artificial intelligence: Commission outlines a European approach to boost investment and set ethical guidelines’.

¹⁶This point has also been made by Thomas Burri (31 May 2018) <www.euractiv.com/section/digital/opinion/the-eu-is-right-to-refuse-legal-personality-for-artificial-intelligence/#comment-334495>.

AI systems should be developed, and whether conferral of legal personality may be an aspect of such an approach.

2.1. Legal personality of AIs

Is there a reason why qualifying an AIs as ‘mere things’ is more puzzling than with any ordinary (nonintelligent) artifact? The intuitive reply is that AI makes autonomous decisions, while an ordinary artifact, no matter how complex, mechanically performs what its human users or designers have predetermined.

AIs can be said to be autonomous in the sense that they are able to supplement and expand their initial epistemic and practical knowledge when it turns out to be incomplete or inadequate for the operating context, and to act upon this extended knowledge, a knowledge not accessible to their users and developers. Autonomy is a graduated property: the greater the readjustment capacity in the face of scarce information, the greater the independence of the artificial agent from the programmer’s instructions and, as a result, the greater the autonomous agency.

Thus, unlike passive artefacts, AIs can be assigned tasks to be performed without human review, thanks to skills like initiating goal-directed behaviour, learning from experience, perceiving the environment and adapting conduct to it, and, if necessary, changing their utility function to maximise the achievement of users’ goals.¹⁷ Autonomous skills frequently correspond to a lack of control on the part of human users, and this is crucial if AIs are involved in legally relevant activities, like entering into a contract or negotiating the terms of an agreement.¹⁸ In fact, since there are multiple courses of action an artificial agent can follow in order to achieve the objective at hand, it is often impossible to predict which of these paths it will actually follow under changing circumstances.

Of course, when artificial agents not having legal personality act autonomously, relying on their own cognitive-behavioural abilities, they don’t bear the ensuing legal consequences (rewards or liabilities). This may lead to an accountability issue: either humans are held accountable for actions they may not have intended to take, nor could reasonably have prevented, or there is a liability gap, as nobody can be held accountable.

On the other hand, we tend to engage with AIs as intentional agents and to assign cognitive-like states to them, so as to explain and predict their behaviour. It is possible to assume that a computer-based artifact acts on the basis of information about one or more states of the world (epistemic states), which it uses to achieve certain goals (conative states) through a rational and flexible decision-making process.¹⁹ Indeed, AIs generally pursue objectives in a way that is not fixed but changes according to their representation of the current state of the external environment and any change in it. Under these circumstances, the conduct of complex AIs can only be interpreted as the output of independent intentional states: only by adopting an ‘intentional stance’ can we predict and explain their behaviour. From the legal standpoint this has consequences on the way

¹⁷Samir Chopra and Laurence F White, *A Legal Theory for Autonomous Artificial Agents* (University of Michigan Press 2011).

¹⁸Giovanni Sartor, ‘Cognitive Automata and the Law: Electronic Contracting and the Intentionality of Software Agents’ (2009) 17 *Artificial Intelligence and Law* 253.

¹⁹Sartor, ‘Cognitive Automata and the Law: Electronic Contracting and the Intentionality of Software Agents’ (n 15) 266.

third parties interact with AIs: they have to rely on the intentional attitudes directly exhibited by the system regardless of the human controllers' upstream intentions. As a result of the properties described so far, it is rather problematic to legally categorise AIs as passive or anyway fully other-directed 'things', even if, as mentioned, one can legitimately disagree about the sufficiency of such operational aptitudes to entirely justify the conferral of legal personality, under the assumption that dignity or conscience, rather than the capacity for goal-driven behaviour, provides grounds for rationality.

However, the mere fact that AI systems possess advanced cognitive capacities may not be a conclusive factor. Certain cognitive capacities – in some regards superior to those of existing AI systems – are in fact also possessed by other entities (e.g., animals), and this does not seem sufficient to justify the conferral of legal personality on such entities, as is also emphasised by the argument from marginal cases. Yet what seem to characterise AIs, to the point of introducing new elements into the debate on legal personality, are the sociotechnical profiles resulting from the deployment of artificial intelligence agents, e.g., the marked unpredictability of their decision-making processes and the impact (both positive and negative) that these processes may have on people's lives, on society, and on the market; the ability of such systems to communicate and network; the involvement of different human players in the production and implementation of such systems, each with different potential responsibilities; and the difficulty, sometimes the impossibility, of tracking the relevant human players. The issues these systems give rise to require a fresh analysis order to determine whether a conferral of legal personality is preferable to other, more limited kinds of empowerment and protection.

2.2. Legal responses and coordination issues

AIs are artefacts, and even if they are very sophisticated, this makes them candidates primarily for the legal category of 'things', since it may seem that artefacts are ultimately other-directed: whatever they are or do is ultimately determined by choices of others, namely, their designers, producers, or users. This status has decisive implications for legal capacity and responsibility, setting aside damages caused by evident programming or manufacturing defects and focusing only on those caused by the *proper* functioning of AIs, or by the interaction between data, inference models, and human parties.

Although the nature of a thing can require a different standard of care, tortious and criminal liability for damages caused by things usually falls on the custodian (who may be the owner or also another controller) or producers (designers-developers). This approach may lead to inadequate legal conclusions in cases in which harmful consequences are caused by an AI system that was nondefective when put on the market, and that was handled with all due care by its users and guardians. Such a system may have engaged in harmful behaviour – e.g., causing a car crash, a loss in online trading – if it has been tricked by the input data, by the limits of its reasoning and learning capacities, as applied to unexpected circumstances. In fact, just as humans are fallible, so are artificial systems, whose learning is based on statistical models. Thus, a system – like a human expert – may commit mistakes even when it meets all state-of-the-art requirements and is deployed to a task that is appropriate to its capacities.

In these situations, conventional liability models may prove inadequate. In case of tortious liability, either a fault is established for negligent supervision by custodians or strict

liability is placed on owners or custodians regardless of any failures of them in controlling the AIs. In fact, in some cases custodians cannot be blamed for negligent supervision, given the autonomy of AIs and the technical inability to foresee their specific behaviour. On the other hand, strict liability risks being unfair, or at least too severe, for the custodian, especially where the damage that needs to be redressed is unexpectedly high. Placing strict liability on the shoulders of producers or programmers, on the other hand, risks discouraging production and innovation. Even greater problems may arise in criminal law, as the cognitive preconditions for criminal liability may not be met by the humans using AI systems.²⁰

A possible solution to these issues, still without attributing legal personality to AIs, may consist in creating an organisation having legal personality – e.g., a limited liability company, a single-member company, or even a memberless LLC (under German law) – which carries out its activity through the AIs.²¹ Alternatively, part of the company's assets can be tied to specific AI-driven businesses (under arts. 2447ff. of the Italian Civil Code). However, presuming that the number of autonomous systems will be increasing in the coming years, the creation of a dedicated company for each autonomous system might be costly, inefficient, and in some cases fiscally demanding.

The aim of setting up a separate patrimony through which creditors can be secured can also be served by conferring a separate legal personality on AIs. This could entitle AIs to hold rights and obligations on their own, enter into contracts by themselves, and produce legal effects on third parties – all this, perhaps, even with complete financial autonomy.

Yet legal personality is a cluster concept: it can result from multiple incidents and it comes with different effects.²² In this perspective, it is necessary to specify the content of this hypothetical 'electronic person' status. But, even assuming a scenario in which there is widespread (political) agreement to assign legal personality to AIs, different sets of rights and obligations will be required depending on the technical peculiarities of AI systems, and maybe the areas in which they are used. It should be also established whether the resulting personality covers all legal areas or just some of them (e.g., only civil law); these are sources of complexity which often prevent us from taking a clear position on the issue, and which disrupt coordination mechanisms when decisions of this kind have to be made in heterogeneous legal contexts.

In other words, there would be a sort of coordination problem, both within and among legal systems. This problem can be broken down into three problems corresponding to three types of asymmetries:

(a1) *Technological*. Each AI can perform very different tasks with various kinds of risk, autonomy, and social impact; the technological peculiarities of each of these tasks

²⁰Francesca Lagioia and Giovanni Sartor, 'AI Systems under Criminal Law: A Legal Analysis and a Regulatory Perspective' (2020) 33 *Philosophy & Technology* 433.

²¹Company law actually provides us with several models under different legal systems, as shown by Shawn Bayern and others, 'Company Law and Autonomous Systems: A Blueprint for Lawyers, Entrepreneurs, and Regulators' (2017) 9 *Hastings Science and Technology Law Journal* 135.

²²Ngairé Naffine, *Law's Meaning of Life, Philosophy, Religion, Darwin and the Legal Person* (Hart Publishing 2009). A slightly different reading of legal personality as a cluster property is provided by Visa AJ Kurki, *A Theory of Legal Personhood* (Oxford Legal Philosophy 2019).

(as concerns both software and mechanics) may call for specifically tailored regulations.

- (a2) *Intra-system*. Different legal fields (e.g., criminal law and company law) may assign slightly different meanings to the notion of personhood, entailing differentiated rules and outcomes; these mismatches often lead to disagreement and misunderstanding among lawyers and legal scholars.
- (a3) *Inter-system*. Each country's legal system may have its own set of rules concerning legal personhood, generally exercising exclusive jurisdiction over these rules; this is a matter of relevance in contexts like that of Europe.

Thus, although there are some functional reasons for granting legal personality to AIs, it would be worth clarifying what use should be made of the concept of personality and how these coordination problems could be solved.

In this paper we do not examine which specific types of AIs 'deserve' to become persons in law, nor do we identify the fundamental reason, the silver bullet, that would in general justify an attribution of personality. We instead argue that jurists should rely on metalegal concepts in managing complexities of this kind, with a view to finding equilibria in the legislative debate; this approach could also work for the legal personality ascribable to AIs. We think that the 'noisy' and uncertain debate on AI personality shows that a reflection on legal personality is needed to address issues pertaining to the conferral and contours of legal personality: some conceptual order is required even before entering into value judgments.

3. Personality statuses and persons in law

Generally, being a person in law means being the holder of legal positions like rights, duties, liberties, powers, and responsibilities. Each position often implies further legal effects for the holder and a chain of jural correlatives for others.²³ More precisely, we can specify as follows the minimal concept of a person in law, i.e., the *thinnest* notion of personality: an entity has personality under the law if there is at least one rule conferring a right or a duty on the entity (or the possibility to acquire a right/duty by triggering the required operational fact). This thinnest concept of legal personality – the abstract possibility of having rights and duties, to which we will refer as *legal subjectivity*²⁴ – can be expanded into different institutional personality statuses, each being conferred on a different kind of entity and having distinct legal consequences.²⁵

A legal system may then have different *thicker* notions of personality, i.e., different personality statuses that, for different kinds of entities, trigger different bundles of rights and duties. A personality status may indeed confer such rights and duties directly or specify the conditions for acquiring them.

²³On correlatives see Wesley N Hohfeld, 'Some Fundamental Legal Conceptions as Applied in Judicial Reasoning' [1913] Yale Law Journal 16.

²⁴A similar point has been raised by Tomasz Pietrzykowski, who recognises an intermediate category of legal entities which he calls nonpersonal subjects of law, in Tomasz Pietrzykowski, 'The Idea of Non-personal Subjects of Law' in Visa A. J. Kurki and Tomasz Pietrzykowski (eds), *Legal Personhood: Animals, Artificial Intelligence and the Unborn* (Springer International Publishing 2017).

²⁵Naffine (n 19) 46.

Typically, in modern legal systems the broadest personality status is conferred on all human beings: they enjoy all fundamental rights and can acquire all kinds of legal positions established by general laws (though some limitation may be established under particular conditions, e.g., incapacity).

More limited personality statuses are conferred on corporations and certain other collective entities whose legal rights are mainly limited to economic relations.

Still more limited personality statuses may be granted to other creatures, such as unborn children and nonhuman animals, and possibly also to natural entities, such as mountains, rivers, and ecosystems. On the interest-based conception of subjective rights, such creatures and entities may be granted legal rights to the extent that the legal system assumes that such creatures and entities have interests of their own that need legal protection (even though the exercise of such rights requires the activity of human agents).

We can also separately consider the different personality statuses that certain entities may have in different areas of the law, such as commercial, family, or criminal law. For instance, we may ask whether corporate entities, while having legal personality under property and tort law, also have personality under criminal law, or whether they have personality under data protection law (having data protection rights).

In conclusion, the function of each personality status can be said to consist in singling out a set of rules (possibly within a specific area of the law) that apply to the corresponding type of entity, so that each instance of that type will be conferred the rights and duties deriving from such rules.

Though there are multiple personality statuses, there is a *central notion* to which we usually refer when speaking of legal personality without further qualifications. This is the general ability to acquire, given generally established triggering conditions, economic rights and obligations, like those arising from contracts or torts. This status is shared by humans and corporations, though humans also possess further legal positions, such as human rights, the ability to enter into family relationships, and their being subject to criminal law (the extent to which such rules may apply to corporate entities is debated and is recognised in certain legal systems only to a limited extent). In the following, ‘personality’ *tout court*, or being ‘a person in law’, will refer to this central notion of a general personality.

4. Personality and the structure of legal kinds

According to the approach of social ontology that we subscribe to, also known as neo-institutionalism, institutions are *systems of rules* that govern social interactions.²⁶

Under these rules, legal kinds – such as contract, marriage, citizenship, and indeed personality – as subspecies of institutional kinds, have a *status function* through which ‘deontic powers’, e.g., rights and obligations – are assigned to objects, persons, entities, or events.²⁷ Which is to say that the thicker notion of personality, presented above, plays a role in linking antecedent facts to legal consequences. So, to understand how

²⁶One of the earliest conceptions of institutions as systems of rules or, more precisely, as ‘rules of the game’, is attributable to Douglass North, ‘Institutions’ (1991) 5 *The Journal of Economic Perspectives* 97.

²⁷John Searle, *The Construction of Social Reality* (Free 1996).

to use the concept of personality, and perhaps legal kinds in general, we need to answer the following questions:

- (a) Under what conditions is an entity considered a person in law (*trigger conditions*)?
- (b) What consequences follow from having personality (*legal implications*)?
- (c) What set of facts explains/justifies why the trigger conditions activate the legal justifications (*background reasons*)?

By contrast, the thinnest notion of personality – namely legal subjectivity, i.e., the unspecified ability to have or acquire at least one right or duty – is not an institutional concept, since it does not specify what rights or duties may be acquired by a legal subject or under what conditions.

At any rate, different theories of legal validity and of the source of legal kinds can arise depending on how background reasons and their connection to positive law are understood.²⁸ While (a) and (b) concern the way legal status works pursuant to the law in force, (c) is constantly being debated among jurists and philosophers, especially when different claims arise as to whether or not certain entities should be qualified as legal persons, as happens with unborn children, nonhuman animals, natural resources, and, in our case, AIs.²⁹ Of course, the first two aspects are a matter for jurisprudential reflection, too: the normative web around legal personhood is in itself a puzzle over which jurists frequently disagree.

We make use of this analytical framework to approach legal personality. From this perspective being a person in law is an institutional-legal fact, positive law rules specify the ‘trigger conditions’ and ‘legal implications’ of personality and further factors provide ‘background reasons’ that support such trigger conditions and legal implications.

Thus, it is important here to disentangle two levels in our analysis of legal personality.³⁰ The ‘trigger conditions’ and ‘legal implications’ set forth in legal rules frame the relation between certain facts and legal properties; e.g., the fact that ‘*x is a human being*’ (1) leads to (a1) the fact that ‘*x is a person in law*’ and therefore also to (b1) all implications associated with general legal personality (rights, duties, power, liabilities, etc.). Thus, the legal status of personality obtains because legal rules exist (e.g., the rules stating that human beings have legal personality, and that personality has certain implications) that establish the trigger condition (a) and the legal implications (b).³¹ This noncausal and synchronic relation of priority is generally called ‘metaphysical grounding’. Under this relation, the more fundamental facts *ground* the less fundamental

²⁸This is Brian Epstein’s insight in his model of metaphysical anchoring relations applied to Hart’s theory: ‘Theories of the criteria of legal validity are theories of the sorts of facts that anchor primary rules. And theories of the sources of legality are theories of the sorts of facts that anchor secondary rules,’ in Brian Epstein, *The Ant Trap* (Oxford University 2015).

²⁹Visa AJ Kurki and Tomasz Pietrzykowski, *Legal Personhood: Animals, Artificial Intelligence and the Unborn* (Springer International Publishing 2017).

³⁰To be more precise, from the point of view of social ontology, legal personality can be understood as a social *property* entailing a relationship between a bearer of rights and duties, the legal system, and third parties. In any case, nothing prevents us from speaking of it in terms of a social *fact*; e.g., ‘Asimov is a person in law in the Japanese legal system.’ It does not in any case change metaphysical relations between facts.

³¹These rules and facts should be considered as complementary in constituting the social property of legal personality.

ones, without being causes of them:³² facts (a1) and implications (b1) obtain *by virtue of* fact (1), and not vice versa.³³

The background reasons (c) for (a) and (b) are a further set of facts, namely, those that support the constitutive rules establishing the grounding relation. These facts can pertain to legal practice – e.g., enactments in conformity with legal procedure, precedent, and judicial decisions³⁴ – or they may pertain to the wider practice that law is part of – e.g., social values and conventions. In the contemporary literature on social ontology there is a debate about the nature of the connection between background reasons and grounding. While some consensus exists on there being a difference between the ‘move-making facts’ grounding different social kinds on the basis of constitutive rules, on the one hand, and the ‘rule-setting facts’ supporting the existence of such rules, on the other, there is controversy over whether the type of metaphysical relationship is the same.³⁵

Brian Epstein claims that there is a peculiar kind of metaphysical link between rule-setting facts and the existence of social rules – a link, not corresponding to grounding, which he calls *anchoring*.³⁶ Anchoring provides *glue* that holds a social kind together through the exposure of the facts that justify its grounding conditions.³⁷ According to Epstein, there are different theories about anchors, and most of them detect the anchoring facts in conventions, shared beliefs, and social practices.³⁸ For instance, on this view, the anchor of a given institutional fact in John Searle’s theory³⁹ is the community’s collective acceptance of the corresponding constitutive rules. So, the ‘anchors’ ensure that there is a grounding connection between priority facts and nonpriority facts (e.g., institutional facts), according to the constitutive rules.

Jonathan Schaffer, on the other hand, insists that for both sets of facts – ‘move-making’ and ‘rule-setting’ facts – the underlying metaphysical connection lies in the relation of grounding, and that, together, the two sets of facts contribute to grounding a social kind. On this view, sometimes labelled ‘conjunctivism’, social kinds are grounded simultaneously in the facts stipulated by a social rule (R) – e.g., a piece of paper issued by the ECB is a euro banknote – and in the facts that cause the social rule to obtain – e.g., collective acceptance of (R).

For our purposes, we need not embrace one model or the other. What is of interest to us is the distinction between two layers in the structure of social (and legal) kinds, i.e., the level of trigger conditions/implications of a social kind and the level the corresponding

³²For an introduction to metaphysical grounding see, Kit Fine ‘Guide to Ground’ in Fabrice Correia and Benjamin Schnieder (eds), *Metaphysical Grounding* (Cambridge University Press 2012).

³³Gideon Rosen, ‘Ground by Law’ (2017) 27 *Philosophical Issues* 279. For examples of grounding outside of the legal domain, see Epstein, *The Ant Trap* (n 25) 56.

³⁴Epstein *The Ant Trap* (n 25) 94.

³⁵The debate between Brian Epstein and Jonathan Schaffer revolves around this distinction. See, specifically, Jonathan Schaffer, ‘Anchoring as Grounding: On Epstein’s The Ant Trap’ (2019) 99 *Philosophy and Phenomenological Research* 749 and Brian Epstein, ‘Anchoring versus Grounding: Reply to Schaffer’ (2019) 99 *Philosophy and Phenomenological Research* 768.

³⁶Mari Mikkola, ‘Grounding and Anchoring: On the Structure of Epstein’s Social Ontology’ [2017] 198.

³⁷Epstein, *The Ant Trap* (n 25) 58.

³⁸In particular, Epstein cites David Hume and John Searle’s theories. Hume believed that facts such as ‘being the owner of a piece of land as first occupant’ were anchored in shared beliefs and expected benefits regarding the rule of first occupancy. Searle, as is known, anchors institutional facts to collective attitudes of acceptance. See Brian Epstein, ‘How Many Kinds of Glue Hold the Social World Together?’ in Mattia Gallotti and John Michael (eds), *Perspectives on Social Ontology and Social Cognition: Studies in the Philosophy of Sociality* (Springer 2014).

³⁹John R Searle, *The Construction of Social Reality* (Free Press 1995).

background reasons. In the following we shall refer to the first as the institutional layer and to the second and the meta-institutional layer.

As further explained in the following sections, on the one hand, the set of trigger conditions (a) and legal implications (b) provide the institutional layer, i.e., the move-making facts of personality,⁴⁰ while the background reasons (c) make up the meta-institutional layer, i.e., the rule-setting facts. Background reasons provide answers to questions like the following: what makes it so that the fact of '*x being a company*' grounds the fact of '*x being a person in law*', or, otherwise stated, what grounds the rule according to which corporations are persons in law?

For the moment, let us note that even if personality – like other legal concepts – is aptly described by its 'trigger conditions' and 'legal implications', to answer the question whether certain entities are suited to acquiring personality in law – whether they *should* be granted personality – we need to investigate the reasons underlying the ascription of personality. The longstanding theoretical debate on the nature of legal personhood is indeed mainly a debate about background reasons, i.e., the anchors (or second-level grounds): some authors believe them to be independent of the legal context, while others view them as mostly systemic and contingent.

5. Theories of legal personhood: legalism vs realism and pluralism vs monism

There are many approaches to legal personality. We will here distinguish them based on two oppositions: that between legalist and realist approaches, and that between pluralist and monist ones. Legalists share the idea that legal personality is basically an artifice through which positive legal systems create a bearer of rights and obligations. They focus on positive law rules specifying the 'trigger conditions' and 'legal implications' of personality and on the specific function it plays in a legal ecosystem. As Lawson points out: 'All that is necessary for the existence of a person is that the lawmaker, be he legislator, judge, or jurist, or even the public at large, should decide to treat it as a subject of rights or other legal relations'.⁴¹

As is known, this way of thinking was most clearly expressed in the theory of Hans Kelsen. According to Kelsen, an entity is a person under the law when this entity is an addressee of legal norms, i.e., when an action by that entity may trigger a sanction, or penalty, against it. An entity's legal personality consists in the set of norms that apply to it. Hence, the core of legal personality lies in 'imputation', as happens in the process consolidating legal effects under an autonomous centre of interest within a legal system. For Kelsen, all legal norms govern human conduct; when a collective entity is personified, the regulated actions of certain individuals (those acting as organs or agents of the collective entity) are ascribed to the collective entity.⁴²

Another distinctly legalistic view is that of H. L. A. Hart, who focused on the linguistic use of legal concepts. Hart thought that nothing essentially corresponds to legal terms, since their core meaning has to be found in the function they perform in legal contexts,

⁴⁰As will be made clear in the following sections, there are good reasons to believe that preconditions and consequences complement the institutional status.

⁴¹Frederik H Lawson, 'The Creative Use of Legal Concepts' (1957) 32 *New York University Law Review* 909.

⁴²Hans Kelsen, *General Theory of Law and State* (Harvard University Press 1945).

rather than in some metaphysical foundation: 'legal words can only be elucidated by considering the conditions under which statements in which they have their characteristic use are true'.⁴³ The meaning of 'person' in law is therefore the result of the linguistic practice that jurists are committed to.

We think that the legalistic view is correct as long as it refers to the grounding relation strictly understood, namely, to the trigger conditions (a) and legal implications (b) of legal personality. Such conditions and implications are indeed established by each legal system. However, an exclusive focus on the legalistic perspective may lead to disregarding the background conditions for ascribing personhood, and thus to an insufficient understanding of the link between personality, on the one hand, and socio-technological opportunities and risks, social attitudes, ethical values, and political ideals, on the other. This limitation of the legalistic view is only partly overcome by broadening the legal perspective so as to include not only the grounding conditions explicitly fixed by legal rules, but also other sources of law, such as legal conventions, shared legal principles, and pragmatic-instrumental reasons pertaining to legally protected interests and values.

Contrary to legalists, realists believe that essential (necessary and sufficient) conditions exist which justify legal personality and that these conditions have universal application, independently of the content of particular legal systems.

Since there are different views on the essential precondition for personality, we have at least four subtypes of realism, each seizing on different properties:⁴⁴

- *Religionist*: Since all human beings are equally sacred, they are worthy of legal protection through the conferral of personality regardless of their cognitive abilities.
- *Naturalists*: What matters for legal personality is that human beings have sentience and share a common destiny. This viewpoint can, of course, be used to promote the personality of nonhuman animals possessing sentience.
- *Rationalists*: Legal personality is suited to those who are cognitively able to exercise the corresponding rights in a rational manner. What matters is the ability to engage in legally relevant interactions and to sensibly assume responsibility for such acts. As Naffine correctly points out, these thinkers tend to restrict legal personhood to a narrow circle of individuals, thus excluding 'marginal' human beings (e.g., very small children or severely mentally impaired adults).⁴⁵
- *Relationalists*: The contours of legal personality are not determined by the possession of specific intrinsic traits or capacities, but rather by membership in a 'community of recognition',⁴⁶ where human entities form their identity and values in part by interacting with nonhuman entities (e.g., the environment and its constituent entities). Thus, for instance, vulnerability or dignity can be seen as relational properties.

⁴³HLA Hart, 'Definition and Theory in Jurisprudence' (1953) 70 *Law. Q. Rev.* 37.

⁴⁴This classification is present in Naffine (n 19) 20.

⁴⁵The philosophical roots of this school of thought trace back to Locke and Kant; many are the modern exponents of this orientation. See, among others, John Chipman Gray, *The Nature and Sources of the Law* (2nd edn, Macmillan 1921); Michael S Moore, *Law and Psychiatry: Rethinking the Relationship* (Cambridge University Press 1984); John Gardner, 'The Mark of Mark of Responsibility' (2003) 23 *Oxford Journal of Legal Studies* 157; SM Solaiman, 'Legal Personality of Robots, Corporations, Idols and Chimpanzees: A Quest for Legitimacy' (2017) 25 *Artificial Intelligence and Law* 155.

⁴⁶See Denis Franco Silva, 'From Human to Person: Detaching Personhood from Human Nature' in Visa AJ Kurki and Tomasz Pietrzykowski (eds), *Legal Personhood: Animals, Artificial Intelligence and the Unborn* (Springer International Publishing 2017).

The realist perspective, regardless of the merit of particular theories falling under this umbrella, disregards the role of positive legal systems in shaping the conditions and implications of rationality. We can view it as expressing natural law perspectives, meant to override different choices made by specific legal systems. Or we may view it as advancing specific theories on the background reasons for ascribing rationality, theories in light of which to critically consider the choices made within specific legal systems.

Under the pluralism vs. monism view we distinguish approaches that accept that different conditions may exist for conferring rationality – with different implications and in light of different background reasons – from approaches on which, on the contrary, a single set of conditions always triggers the same implications, and on the same rationale. Obviously, intermediate positions may also exist.

Legalistic approaches are generally pluralistic, as they recognise whatever different grounds for personality may be recognised by different present, past, or even future legal systems.

Realist approaches, on the contrary, skew toward monism, usually focusing on a single ground or rationale for personhood. Furthermore, in terms of the metaphysical structure of legal personality set out in the previous section, realists are more inclined to equate the content of anchors and grounds.

We believe that a pluralistic approach should be adopted, not only at the level of the grounding relation but also at the level of the background reasons (which anchor or ground in a structural fashion). In fact, different justifications may support the conferral of personality of different types of entities, and in different areas of the law.

Thus, for instance, criminal law emphasises cognitive attitudes and rational behaviour, while medical law appears to be influenced by religious or moral values connected with the sacredness of life (this is what emerges, for example, when we turn to issues like the end of life, abortion, and the legal subjectivity of the unborn).⁴⁷ As a result, there may be good reasons to consider an entity a legal person under medical law, but not under criminal law. Finally, as previously emphasised with the ‘argument from marginal cases’, even the criterion of rationality does not provide a univocal guide for attributing moral or, in the case at hand, legal statuses.

Legal personality is therefore not a unitary concept: whether an entity is eligible for personhood and, if so, what kind of personhood it is eligible for, is a question that needs to be contextualised.⁴⁸ The need for contextualisation, however, does not debunk the idea that the ascription of such personality requires reasonable background reasons. Such reasons may be linked primarily to economic purposes, pertaining to the efficient management of resources and activities, as is the case for collective entities, but even more so for impersonal legal platforms like single-member or memberless companies.⁴⁹

6. The nature of legal concepts I: inferentialism

Law demands that we select facts from social practice and then reformulate them according to its own internal patterns. Thus, it creates an independent store of knowledge, so

⁴⁷Naffine (n 19) 164.

⁴⁸Richard Tur, *The ‘Person’ in Law* (Basil Blackwell 1987).

⁴⁹Here the distinction proposed by Visa Kurki between ‘legal persons’ and ‘legal platforms’ seems also relevant. In Visa AJ Kurki, *A Theory of Legal Personhood* (n 19) 133.

that the legal community can effectively coordinate both in implementing rules and in resolving disputes.⁵⁰ This function of the law fits legalist assumptions: legal concepts are constituted by the rules that establish their use conditions and their deontic consequences. In view of this, it is possible to conceive of legal concepts – and so also of personality – as mere links between factual preconditions and normative effects. They are basically ‘means of presentation’ of multiple rules which don’t carry any independent meaning from the link they hold, as Alf Ross believed was the case with all ‘intermediate legal terms’.⁵¹

This standpoint is consistent with a form of semantic inferentialism. In short, rather than identifying legal concepts by their referential function, we can do so by their inferential role: by the role they play as antecedents or consequents in legal reasoning.⁵² On this view, the content of a concept consists of the bundle of inferences endorsed in its practical use and context. Also, to employ a concept is tantamount to being committed to the inferential relations it involves with other concepts – relations such as exclusion, derivation, and implication. From an inferential perspective, legal personality links a set of triggering conditions – e.g., being human – with a bundle of legal norms – e.g., the right to sue and the liability to be sued. As a consequence, legal personality can be viewed as an implicit condition in every rule conferring rights or obligations. For instance, once the personality presupposition is made explicit in:

- If x wrongfully harms y , then x has an obligation to make y whole.

the rule should read as

- If x is a person, and y is a person, and if x wrongfully harms y , then x has an obligation to make y whole.

Inferentialism about legal concepts is consistent both with the legalistic view of personality as an abstract label that makes scattered norms homogeneous and with the idea that legal concepts embed information functional to the practical cognition of social agents. Here there are several relevant implications for the representation of legal knowledge we can benefit from.

First, identifying the meaning of legal concepts with the inferential network can be an effective way to portray the articulation of a cluster concept having disjunctive preconditions and multiple legal effects, like legal personality. Second, this vision emphasises that legal concepts are meant to provide the inferences a particular legal system makes viable. Third, since this makes it possible to understand a general legal term only through the inferential network in which it is incorporated, it is possible to gain an understanding of the different ways in the same term is used in different areas of the law. And, finally, by viewing conceptual links as defeasible inferences, inferentialism can enable a flexible reading of legal concepts, i.e., where the ascription and implications of such concepts are subject to exceptions.

⁵⁰Mariano Croce, *Self-Sufficiency of Law: A Critical-Institutional Theory of Social Order* (Springer 2012).

⁵¹See Alf Ross, ‘Tü-tü’ (1957) 70 *Harvard Law Review* 812; Alf Ross, ‘Definition in Legal Language’ (1960) 25 *The Journal of Symbolic Logic* 90.

⁵²For an introduction to inferentialism Wilfrid Sellars, ‘Inference and Meaning’ (1953) 62 *Mind* 313; Robert Brandom, *Making It Explicit: Reasoning, Representing, and Discursive Commitment* (Harvard University Press 1994) and Robert Brandom, *Articulating Reasons: An Introduction to Inferentialism* (Harvard University Press 2009).

An approach based on inferentialism may help us to figure out the legal consequences of viewing AI systems (or other nonhuman entities) as having general legal personality in the context of specific legal systems, and to find appropriate adaptations by fixing the inferential links involved as needed.

However, the idea that legal personality is an inferential node captures (a) its ‘trigger conditions’ and (b) its ‘legal implications’, but fails to include (c) the ‘background reasons’ that are part of a comprehensive theory of legal personality.

From a merely inferentialist perspective, legal concepts do not have any intrinsic dynamism: their constitutive elements are fixed in by the inferential links established by positive law. But that is not how legal practice actually works. A coevolutionary dynamic is at play between social and legal practice: on one hand, legal concepts are permeable to social facts – whether moral, political, or economic – and constantly adjust to their changes⁵³; on the other, the law provides us with reasons for action, incentives, and motivations, giving rise to extra-legal phenomena. Defeasible inferences that originate from disjointed events and give rise to various legal effects are held together by porous conceptual categories, not fully reducible to legal rules, which also reflect nonlegal knowledge. This wider conceptual framework enables legal inferences to be applied and interpreted contiguously with the social and moral landscape in which legal practice is embedded.

This socioempirical context places strong constraints on how rules are formed and how social agents make use of them. Figuring out how this happens is tantamount to understanding what the ‘background reasons’ of legal personality are and what role they play in shaping specific legal prerogatives.

To conclude, it would be preferable to subscribe to a less demanding version of inferentialism: even though legal concepts are constituted by their inferential role – by the set of rules authorising their inferential use – this would not be sufficient to account for their axiological or teleological fitness.

7. The nature of legal concepts II: a neoinstitutional account of legal personality

To move beyond inferentialism, strictly understood, we will refer to the theory of institutional legal concepts developed by Neil MacCormick. Drawing on Searle’s social ontology,⁵⁴ MacCormick views legal facts as institutional facts.⁵⁵

An institutional fact is conceived as something whose existence depends ‘not merely upon the occurrence of acts or events in the world’ – as is the case with brute facts – but also upon rules.⁵⁶ Indeed, institutions, from which institutional facts descend, presuppose a system of rules of a peculiar kind: constitutive rules. In a nutshell, constitutive rules generally have the logical form ‘X counts as Y in context C’, where Y stands for the *status function* these rules assign to people or objects within a given context C.⁵⁷

⁵³Simon Deakin, ‘Juridical Ontology: The Evolution of Legal Form’ (2015) 40 *Historical Social Research* 170.

⁵⁴Searle (n 36).

⁵⁵Neil MacCormick, ‘Law as Institutional Fact’ in Neil MacCormick and Ota Weinberger (eds), *An Institutional Theory of Law. Law and Philosophy Library* (Springer 1986) 51.

⁵⁶This distinction between brute and institutional facts can originally be traced back to Gertrude EM Anscombe, ‘Modern Moral Philosophy’ (1958) 33 *Philosophy* 1.

⁵⁷John Searle, *Making the Social World: The Structure of Human Civilization* (Oxford University Press 2009).

The assignment of status functions, Searle claims, cannot be explained by the physical makeup of those people or objects but is rather the result of collective intentionality, such as our disposition to cooperate and share desires and beliefs about the institutional condition of certain entities. The status so recognised triggers further consequences, which Searle calls ‘deontic powers’: e.g., rights, obligations, requirements, permissions, authorizations, entitlements. Money, for instance, is an institution insofar as it presupposes a system of constitutive rules assigning an agreed exchange value to certain objects (X will *count as* money in C); defined functions will correspond to those objects and a series of powers and duties will be entrusted to those who use them.

Against this background, MacCormick maintains that systemic legal concepts – such as property, marriage, succession, and personality – are institutions of law whose specific instances are institutional facts. For instance, the legal personality of a company is an institutional fact that relies on the institution of legal personality. Such concepts, which he calls *institutional legal concepts*, are systemic in that they hold together many interrelated legal rules. They are used, in fact, as tools for sorting out intricate normative content.

According to MacCormick, each institutional legal concept is regulated by a triad of rules:

- (1) *Institutive rules*: Rules linking the occurrence of certain acts or events to the coming into existence of a specific instance of a legal institution. Consider, for instance, the rule stating that if a group of people allocate money to a common asset for an associative purpose, then an association having legal personality comes into existence.⁵⁸
- (2) *Consequential rules*: Rules determining the implications of the existence of each specific instance of an institution; e.g., if a legal person exists, then this person can be held responsible for any loss, damage, or injury it causes.
- (3) *Terminative rules*: Rules laying down the conditions under which a specific instance of an institution of law ceases to exist and thus ceases to produce legal effects; e.g., the legal personality of a collective body dissolves when the purpose of the association becomes unattainable.

Searle’s and MacCormick’s positions are compatible insofar as we accept the existence of two types of constitutive rules: *by condition* and *by implication*.⁵⁹ Basically, this means that institutional concepts – like money and borders – but also legal ones – like personality, contract, and ownership – are defined not only by their conditions of existence but also by their consequences. In this way, MacCormick’s institutive and terminative rules would appear as constitutive *by condition*, while his consequential ones as constitutive *by implication*.

At any rate, legal personality can be conceived as an institutional legal concept: a systemic concept made up of institutive, consequential, and terminative rules. This way of conceiving legal personality yields an account of the temporal persistence of its specific instances, which come into existence as a result of certain acts or events, operate through

⁵⁸Dick WP Ruiters, *Institutional Legal Facts: Legal Powers and their Effects* (Kluwer Academic Publishers 1993).

⁵⁹This idea was suggested by Corrado Rovorsi under the name of ‘complementarity thesis’ in Corrado Rovorsi, *Costituire: Uno studio di ontologia giuridica* (Giappichelli Editore 2012). See also, Frank Hindriks, *Rules & Institutions: Essays in Meaning, Speech Acts and Social Ontology* (Haveka B. V. 2005).

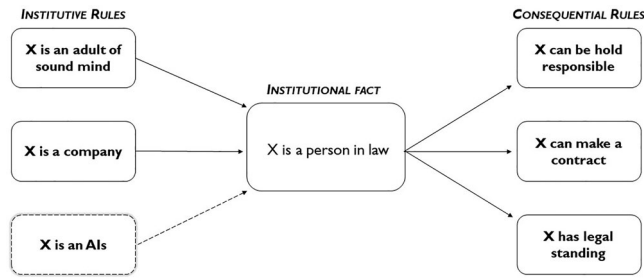


Figure 1. Inferentialism and MacCormick's institutionalism on legal personality.

sets of legal implications like rights, duties, powers, and responsibilities ('status functions' and 'deontic powers'),⁶⁰ and finally end at a given moment (Figure 1).

MacCormick's view is compatible with an inferential reading of legal concepts since the institutionalist architecture looks consistent with the legalist idea that personality is a label through which the legal system organises fragmentary rules in order to channel them toward a single point of imputation. However, MacCormick tries to reconcile the legalistic view and the social dimension of legal phenomena. In fact, MacCormick claims constitutive rules alone cannot provide an adequate account of institutions.⁶¹ He argues that a profitable understanding of institutions would instead make it necessary to lay bare their 'underlying principle' or 'final cause'. Hence, an institution and its specific instances can only be understood in light of the general purpose they serve, in the context within which they operate. As MacCormick exemplifies, 'corporations are associations of individuals to which a separate legal personality attaches for the purpose of holding property and bearing and discharging legal obligations and responsibilities'.⁶²

From MacCormick's perspective a legal institution can be approached by bringing three distinct things into focus. Firstly, we may focus on the institution itself as the set of constitutive rules. Secondly, we may focus on single instances of the institution, i.e., institutional legal facts, that are created by concrete behaviours and events appropriately matching constitutive requirements. Thirdly, we may focus on the social practice supporting the institution, namely, on the social and individual interests the institution is meant to satisfy, and on the individual and collective action aimed at such interests.⁶³

In the next section we shall further develop MacCormick's perspective by exploring the possibility of integrating these aspects (in particular the third one, which goes beyond the usual account of institutional facts), to which end we will be relying on the path charted by social ontology.

8. Developing the neoinstitutional ontology: meta-institutional concepts

One of the most cited arguments in favour a broad understanding of social ontology – not limited to constitutive rules – comes from the thought experiment proposed by

⁶⁰Searle's terminology in (n 36) and (n 54).

⁶¹Neil MacCormick, 'Norms, Institutions and Institutional Facts' (1998) 17 *Law and Philosophy* 301

⁶²*ibid* 314.

⁶³Neil MacCormick, 'Persons as Institutional Facts' in Ota Weinberger and Werner Krawietz (eds), *Reine Rechtslehre im Spiegel ihrer Fortsetzer und Kritiker* (Springer 1988).

Hubert Schwyzer, who highlighted that there is something behind the constitutive rules of an institutional practice.⁶⁴ Schwyzer imagines a society – the ‘Ruritanians’ – in which the game of chess exists with its typical rules, but instead of being a competitive game it is a religious ritual. It follows that Ruritanian chess lacks the concept of ‘victory’ or ‘defeat’ and that the actions associated with such concepts cannot be carried out. Yet the concepts of victory and defeat are logically independent of the constitutive rules of chess, since they are rooted in the wider social practice in which the game of chess is immersed, i.e., the practice of competitive game playing. These concepts would lie in the background of constitutive rules, a further level composed of concepts that are presupposed by institutional practices. That is, there are factors which are not specified by the rules of a social practice, but which bring the institutional practices into being and, at the same time, enable the practice to function in its distinctive way. These factors are covered by *meta-institutional concepts*.⁶⁵

Back to our topic, let us see what further levels can contribute to more adequately framing and explaining legal personality, thereby ultimately giving us a better handle on the issue of AIs. According to the inferentialist account it is possible to explicate the concept of legal personality for natural persons in this way:

- (1) If z is a human being, then z has legal personality (*constitutive by condition*).
- (2) If z has legal personality, then z can enter into contracts (*constitutive by implication*).

Rules (1) and (2), when applied to a specific fact (e.g., Jane Doe is a human being) converge into the same institutional fact: ‘Jane Doe has legal personality’, which entails that ‘Jane Doe can enter into contracts’. In MacCormick’s terms, this fact (‘Jane Doe has legal personality’) counts as a specific instance of an institution of law. But this institutional fact also seems to depend on presuppositions that are not defined through the set of constitutive rules of legal personality. In the case here proposed, (2) seems to assume the capacity of the entities at stake (those being granted legal personality) to act in accordance with their rights and duties, to intelligibly communicate, and to rationally deliberate. The family of presuppositions of legal personality is not a fixed set but rather contains social justifications and moral perspectives that may change over time. For example, where legal personality manifests as the possibility of acquiring rights and obligations, or to sue and be sued, agency and moral competence are taken for granted. In other cases, where, for example, legal personality manifests merely as instrumental to the protection of interests (i.e., *passive* legal personality),⁶⁶ values or facts – e.g., vulnerability or integrity – are presupposed relative to the legitimacy of the protection provided by the status. In any case, these presuppositions are captured by concepts that appear to be *prior* to that of legal personality, in the sense that they condition the possibility of the institution itself.⁶⁷ For instance, the concept of agency does not originate with the practice of law itself but rather emerges from a wider context

⁶⁴Hubert Schwyzer, ‘Rules and Practices’ (1969) 78 *The Philosophical Review* 451.

⁶⁵The expression ‘meta-institutional concepts’ was first introduced by Dolores Miller, ‘Constitutive Rules and Essential Rules’ (1981) 39 *Philosophical Studies* 183.

⁶⁶MacCormick, ‘Persons as Institutional Facts’ (n 60) 383.

⁶⁷Giuseppe Lorini, ‘Meta-Institutional Concepts: A New Category for Social Ontology’ (2014) 56 *Rivista di estetica* 127.

which law is part of. To some extent, it is precisely the role of law in a broader social practice that makes it so that agency should be presupposed.⁶⁸

Meta-institutional concepts linked to legal personality are not constituted by the normative conditions for being a ‘person’ in law, but rather amount to moral, social, or political factors. Yet it would not be possible to design the constitutive rules of legal personality or to critically reflect upon such rules without implicitly relying on some of these notions. Depending on the concept (or family of concepts) presupposed, differentiated personality statuses may be obtained.⁶⁹

What does the meta-institutional level consist of? It cannot be made up of institutional rules having the form ‘*x* counts as *y* in *C*’, since in this case meta-institutional concepts would not be distinct from institutional concepts. It also cannot be made up of mere brute facts, since in this case it would be incapable of providing normative justifications/explanations for institutions. Rather, it is made up of social values, conventions, standards of conduct, and shared beliefs which are (at least partly) exogenous to the institutional rules. It includes different kinds of concepts: some of them are *axiological* (e.g., ‘justice’ in the legal domain), while others are mainly *teleological* (e.g., ‘victory’ in chess).⁷⁰

Meta-institutional concepts transcend the boundaries of single institutions – hence the boundaries of the structure set up with constitutive rules – and give meaning to the institutional practice within its socioempirical environment. In the case of personality, such concepts – ideas of humanity, agency, sentience, integrity, environmental value, economic expediency, etc. – reflect individual and collective attitudes in different fields, values that both drive legal practice and provide a common ground between the legal, moral, and sociopolitical domains. The institution of legal personality can truly be ‘played out’ only if contextualised within this background that meta-institutional concepts display.

9. Theoretical benefits of a multilayered ontology

In the previous sections, the inferential reading of institutional concepts was found to be consistent with an ontological apparatus that also includes the backstage of meta-institutional concepts. Now, what advantages derive from this joint framework? In pointing these out, it will be useful to distinguish the purely theoretical implications from the practical ones (especially from the focus on the legal status of AI).

The theoretical advantages – i.e., for a general understanding of legal concepts – are multiple. In the first place, the meta-institutional level shows that an institution’s constitutive rules do not emerge out of nothing but take shape within a conceptual and ‘semantical atmosphere’.⁷¹ The teleological and axiological landscape influences institutional practices, e.g., concepts of ‘moral agency’ or ‘moral patience’ imposes some constraints on legal personality itself. It is thus possible to highlight some of the beliefs the legislator

⁶⁸Corrado Roversi, ‘Conceptualizing Institutions’ (2014) 13 *Phenomenology and the Cognitive Sciences* 201.

⁶⁹Relevant to the pluralism of determinants of personhood – both legal and moral – is Gellers’s multi-spectral approach in Gellers (n 8) 151.

⁷⁰Roversi, ‘Conceptualizing Institutions’ (n 65) 205.

⁷¹Giuseppe Lorini and Wojciech Żelaniec, ‘The Background of Constitutive Rules: Introduction’ (2018) 4 *Argumenta* (Special Issue), 13.

is committed to when making decisions about legal policy with respect to who should be recognised as having subjective legal positions and how. This point reinforces MacCormick's claim about the inadequacy of an account of institutional practices solely based on constitutive rules: 'unless you know the underlying principle or final cause of a given institution, it profits you nothing to know how an instance of it can be established'.⁷²

Second, a multilayered ontology shows that legal institutions are not self-contained, but rather that some external concepts are presupposed. An awareness of these interlinkages can play a significant role in legal cognition itself.⁷³ This reinforces the idea that legal forms tend to be constantly adjusted to social practices.⁷⁴ In addition, this multilayered ontology provides conceptual gateways between the different degrees of existence of institutions in MacCormick's theory (namely, between the social dimension and the juridical one).

Third, a multilayered ontology enables us to understand the merit of having cluster concepts, such as legal personality, covering, or potentially covering, disparate situations (e.g., humans, children, companies, animals, AI systems). In fact, the meta-institutional level brings into account interests and values – which evolve depending on the period and the social system in question – grounding distant circumstances for the possible ascription of personality. Occasionally, this background may highlight affinities or distinctions, enabling or disabling analogical connections. As a consequence, the coexistence of *thick* and *thin* notions of legal personality also becomes comprehensible.

Fourth, a theory of legal personality should address not only the 'trigger conditions' and 'normative implications' of such a legal status (see *supra*, Section 3), but also its 'background reasons', that is to say, what facts bring into play the rules framing the conditions of legal personhood. We cannot here discuss pluralism or individualism about grounds or anchors⁷⁵; what matters for us is that meta-institutional concepts incorporate background reasons highlighting the deeper facts and values that are presupposed by legal institutions and their constitutive rules; that is to say, they describe so-called 'rule-setting facts'. On this model, the metaphysical structure of legal personality bears at least two different relationships: a *structuring* grounding – to quote Schaffer – between facts at the meta-institutional level and the constitutive rules of legal personality; and a *triggering* grounding between move-making facts and the institutional fact/property of legal personality.⁷⁶

Detecting structuring grounds (or anchors) is useful for a better justification of rules in legal enforcement. They point out prior general values and social conventions able to shape our choices in the legal sphere. Even though legal concepts are constituted by inferential links, if we are to justify particular inferences and then commit to their applicability, we need these concepts and inferences to be consistent with some deeper reasons. Law is able to provide reasons for action and motivate individuals thanks also to the involvement of deeper conventions rooted in the social practice in which the law is embedded.⁷⁷

⁷²MacCormick, 'Norms, Institutions and Institutional Facts' (n 55) 319.

⁷³Sartor, 'Cognitive Automata and the Law: Electronic Contracting and the Intentionality of Software Agents' (n 15) 218.

⁷⁴Deakin (n 49) 182.

⁷⁵For an interesting reading of pluralism and individualism about grounds and anchors – in partial disagreement with Epstein – see Francesco Guala, 'Epstein on Anchors and Grounds' (2016) 2 *Journal of Social Ontology* 135.

⁷⁶Schaffer (n 32) 755.

⁷⁷On deep conventions see Andrei Marmor, 'Deep Conventions' (2007) 74 *Philosophy and Phenomenological Research* 586.

Finally, closely related to the previous point, meta-institutional concepts can be employed in the legal field to test and reformulate the content of institutional concepts.⁷⁸ In the case of personality, such concepts can help us question who or what ought to have the legal protections and empowerments linked to personality and in which way.

This last point bears some elaboration. Sometimes this review is done indirectly: rather than referring explicitly to meta-institutional notions, auxiliary concepts are used to bridge the indeterminacy of meta-level pro and con reasons with the (relative) determinacy of legal rules. The function of these auxiliary concepts, then, is similar to that which in moral philosophy is served by what are referred to as *midlevel principles*, which help to justify rules and particular judgments by specifying or balancing higher-level principles.⁷⁹ By going back and forth between deep reasons and the particular assessments, we make in concrete cases is what characterises the method of *reflective equilibrium*.⁸⁰ Through this method, coherence is sought ‘among the widest set of moral and nonmoral beliefs by revising and refining them at all levels’.⁸¹ As we shall see in the next section, auxiliary terms can also be used to achieve a kind of equilibrium in the legislative debate.

In the area of legal personhood, *legal subjectivity* may be viewed as an auxiliary transitional concept of this kind, coinciding with the thinnest notion of personality (having, or having the ability to acquire, some legal rights or duties). This concept is meant to capture, at least in Italian legal commentary and jurisprudential debate, potential eligibility for personality status. Legal subjectivity is thus mostly used to underpin some entities’ merits or competence in relation to legal rights, frequently in a mitigated or conditioned manner (consider, for instance, condominiums or unborn children). However, legal subjectivity is not an institutional concept proper, since it does not refer to a precise set of legal rights. With this notion, we can proceed on either a descriptive stance or a programmatic one: in the former case we point to the minimal personality status, i.e., we point out that some entities are already recipients of at least one legal position (*thinner* notion); in the latter case we claim that these entities should acquire some broader personality status.

10. Applying a meta-institutional and an institutional perspective to the case study

The approach we have introduced, which distinguishes meta-institutional and institutional aspects, can be useful for examining the question of the legal personhood ascribable to AIs, since it enables us to distinguish criteria of ascription at different levels in different contexts.

A meta-institutional perspective enables us to ‘test’ whether an entity possesses the general properties and attitudes that may justify an ascription of personality to it: cognitive capacity, vulnerability, dignity, the ability to have interests, moral patience, sentience,

⁷⁸Roversi, ‘Conceptualizing Institutions’ (n 65) 208.

⁷⁹Michael D Bayles, ‘Mid-level Principles and Justification’ in J Roland Pennock and John W Chapman (eds), *Nomos XXVIII: Justification* (New York University Press 1986).

⁸⁰On ‘reflective equilibrium’ see John Rawls, *A Theory of Justice* (Harvard University Press 1971) and Norman Daniels, *Justice and Justification: Reflective Equilibrium in Theory and Practice* (Cambridge University Press 1996).

⁸¹Daniels (n 77) *ibid* 2.

social role, the ability to communicate, economic expediency, etc. Cognitive capacity, for example, is one element supporting an ascription of personality to AIs, and another supporting element could lie in their ability to facilitate certain social or economic interactions.

However, none of the properties identified at the meta-institutional level is sufficient, separately considered, to ground legal personality, i.e., to determine whether such entities should be granted legal personality. To this end, we need to balance the advantages and disadvantages that would obtain if legal personhood – as a general ability to have rights and duties, at least in the patrimonial domain – were to be conferred on such entities. Consider, for instance, the multiple implications involved in abortion, inheritance, medical liability, etc., that could result once legal personhood is conferred on the unborn child. This judgement may be facilitated if we preliminarily test the candidates for legal personality by resorting to intermediate concepts, such as the idea of legal subjectivity. This midlevel review may also consist of a judgment of expediency, which may also eventuate in the claim to personhood being rejected. It may be concluded that certain entities we have classified as legal subjects may not require personality, since the protections, guarantees, or enabling conditions that are suited to such entities (e.g., unborn fetuses, animals, ecosystems, technological systems) may already be secured under different legal regimes that are better suited to such entities. For instance, it may be argued that sentient beings like nonhuman animals should not be qualified as legal persons, since they lack the cognitive capacity to understand their legal positions and act accordingly. On the other hand, it may be argued that cognitively capable beings, such as advanced AI systems, do not deserve the protection that is granted to legal persons, since they lack sentience, being unable to experience proper feelings. Or, from the same midlevel review itself, a judgement of expediency may point in the opposite direction, suggesting, for example, that there may be rhetorical and normative value in conferring legal personality even where certain legal protection are already in place.

On the contrary, let us assume that we conclude that a category of entities should be granted legal personality, with an accompanying set of rights and duties (or with an opportunity to acquire them given the appropriate operative facts). In this case, we need to determine how we should configure legal personality of such entities – with what restrictions or extensions relative to the default idea of personality as the general ability to have rights and duties the patrimonial domain – in keeping with an adequate balance and adjustment of presupposed meta-institutional values. On this basis, an argument can be made to the effect that such entities can already be viewed as legal persons based on existing law (*de lege lata*) or that a change in the law is needed for personality to be conferred (*de lege ferenda*).

With regard to AI systems, the second approach (a change in the law) seems more plausible. It seems inappropriate to grant general legal personality to AI systems merely through legal interpretation, given the novelty of such entities and how different they are from those that have so far been granted legal personality (which differences make analogies highly questionable), and the important political implications of choices about the role that AI systems should play in society.

As long as legislators have not made that choice, the law may account of the existence of autonomous AI systems by relying on transitional concepts like that of legal subjectivity, which conveys the idea that certain entities, given their special features requires some

legal recognition of their interests or capacities. The recognition of the legal subjectivity of autonomous AI systems may be used to build consensus around the need for certain kinds of AI systems to have a legal regime that takes their cognitive capacities, social functions, and accountability gaps into account, granting them some legal protections, or enabling them to autonomously and actively engage in certain legal activities, or just limiting their users' liabilities. Such a regime could be introduced even as uncertainty or even denial persists concerning the attribution of personality.

The case of the unborn child has indeed been addressed in a similar way in Italy. There is no rule in the Italian Civil Code granting legal personality to the unborn, and a permissive abortion law is in force. However, by referring to certain rules which protect human life, in the Italian Constitution and international law, the Italian judiciary has argued that damages are owed for injury to the unborn child. According to the Court, while the unborn child is not a (natural) person, he or she is a legal subject whose interests are to some extent taken into account by the law. The same holds for unincorporated partnerships or associations, which in different legal systems are accorded different levels of legal protection and different legal powers according to their purpose and composition. Interestingly, the same kinds of partnership or association may be conferred legal personality in one legal system but not in another.

As mentioned in Section 2.2, the decision to grant personhood to AI systems can be affected by coordination issues. It is difficult to imagine that a single criterion could ever be used in Europe to determine whether an AI system should count as a legal person.

Indeed, for one thing, the technological heterogeneity among AI systems and the context of their use makes it impossible to resort to a single criterion to determine in what cases such entities could have personality. The conditions that make legal personality appropriate in one context (e.g., e-commerce) may be very different from those that make it useful in another (e.g., robots used in health care or in manufacturing). And, for another thing, we need to consider that in Europe different conceptions of legal personhood exist, both within the same legal system and between different legal systems.

This does not exclude, however, the possibility of the European legislator recognising a particular legal status (or different such statuses) for AI systems which satisfy certain conditions: the fulfilment of technical standards – e.g., transparency, explicability, safety, reliability – as well as certain levels of performance and autonomy, as defined by technical capacities and users' choices. In this sense, it is desirable for the legislator to take a multifactorial approach in which multiple sociotechnical aspects combine to support the recognition of a special legal status for advanced AI systems, albeit in a gradual manner and without necessarily triggering full legal personality. Such a status may come into shape when the users and owners of certain AI systems are partly shielded from liability (through liability caps, for instance) and when the contractual activities undertaken by AI systems are recognised as having legal effect (though such effects may ultimately concern the legal rights and duties of owners/users), making it possible to view these systems as quasi-holders of corresponding legal positions. The fact that certain AI systems are recognised by the law as loci of interests and activities may support arguments to the effect that – through analogy or legislative reform – other AI entities should (or should not) be viewed in the same way.

Should it be the case that, given certain conditions (such as compliance with contractual terms and no fraud), the liability of users and owners – both for harm caused by AI

systems of a certain kind and for contractual obligations incurred through the use of such systems – is limited to the resources they have committed to the AI systems at issue, we might conclude that the transition from legal subjectivity to full legal personality is being accomplished.

The combination of the meta-institutional and institutional perspectives – incorporating social and technological aspects, and supported by a flexible legal subjectivity – can mitigate the problem of coordination we took as a premise: (a1) the meta-institutional level identifies the area of potential ‘electronic persons’, and may be instantiated as a sort of testing procedure; (a2) the auxiliary concept of legal subjectivity would not commit national lawmakers to a specific conception of personality, and they would remain free to carve out the institutional status in detail in accordance with their legal systems; (a3) finally, this would also give legislators, national and supranational alike, the ability to determine whether a personality status will exist under civil, contract, or criminal law, or a combination of the three.

To conclude, the analysis of the background of institutional legal concepts makes it possible to discover notions and theoretical tools that can be effectively implemented, as well as to support the conferral of thick legal status, such as general personality. Auxiliary concepts like legal subjectivity may help us to locate and specify the aspects that are relevant to the legal status under consideration, while also making our set of moral intuitions, social beliefs, and conventions more stable for the purpose of their legal use. The mutual alignment between specific institutional rules and broader meta-institutional values contributes to the formation of an equilibrium among different conceptions of subjecthood and personality across different legal systems as well as across different areas within each legal system.⁸²

Acknowledgments

Giovanni Sartor has been supported by the H2020 ERC-Advanced Project “CompuLaw” (G.A. 833647).

Disclosure statement

No potential conflict of interest was reported by the author(s).

⁸²On the role of midlevel principles for wide reflective equilibrium, see Josep J Moreso and Chiara Valentini, ‘Judicial Dialogue as Wide Reflective Equilibrium: In the Region of Middle Axioms’ (IVR Lucerne 2019).