Knowing the Unknowable: Soft Laws and Hard Decisions

Jeremy Pitt
Imperial College London, London SW7 2BT, U.K.

■ **THE 2021 NATIONS LEAGUE CUP FINAL** was a soccer/football match played between France and Spain. It was settled by a goal scored by a player who was in an offside position when a ball was passed toward him by his teammate, but was not penalized because in its trajectory, the ball scraped the studs of an opponent trying to intercept the ball. According to FIFA's laws of the game "A player in an offside position receiving the ball from an opponent who deliberately plays the ball, … is not considered to have gained an advantage" [1, p. 100] and therefore has not caused an offside offense. The goal was allowed to stand, and France duly won the match, and the cup, with this the winning goal.

This was the cause of some controversy.

It certainly seems bizarre. If his opponent had not touched the ball at all, or if it had been accidentally deflected by his opponent, then the goal-scorer would have been considered to have been in an offside position and would have been penalized for an offside offense. But because the opposing player tried, and failed, to intercept a ball being passed to the player in an offside position (which of course he did not, and could not, know), and in doing so got the merest touch on the ball and this touch was deemed deliberate, it was enough for Schrödinger's footballer to flip between causing, and not causing, an offside offense.

Talk about entrapment.

There was some commentary after the match about the decision being technically correct according to a literal interpretation of the laws of the game, but not being in "the spirit of the game." It is easy to imagine that this situation was not of the type that the regulators had in mind when they drafted the rules, and added the caveat quoted above. But, if this is the value that is being promoted, that is, if the decisions that are produced are judged in consistent with the "spirit of the game" (since no one has any idea what a game spirit looks like or knows what incantation will invoke one to whisper in the referee's ear or act as a witness in the Court of Arbitration in Sport), then the questions become: 1) what is wrong with the laws as they are currently formulated; 2) how perhaps could they be alternately formulated; and 3) under this alternate formulation, what procedure could be used to produce a consistent judgment. And finally, what are the implications for this treatment of "soft laws" in sports arbitration in other domains, for example, in criminal law or deliberative assemblies.

In considering this first question, there are three initial observations. The first is that the laws themselves have not quite gone "The Full Nomic" [2], in the sense that the rules do not specify their own amendment and degenerate into paradox, but they have become increasingly complex and generally resistant to traditional restaurant-tabletop explanations with salt-cellars and pepper-pots. With that complexity, there is a scope for something on the spectrum between ambiguity and inconsistency: in the situation described above, a moderate lawyer

could surely make a case for offside on the basis that "a player in an offside position is … penalized on becoming involved in active play by … clearly attempting to play a ball which is close when this action impacts on an opponent" [1, pp. 99–100].

The second observation is that, as a consequence of that complexity, the introduction of technology and video assistant referee (VAR) has not done anything to diminish the potential for controversy, all too predictably [3]. While the use of technology in sporting adjudication has been beneficial in binary situations and for decision support, for example, all that VAR has contributed (besides delays, expense, and elitism) is an extremely literal and reductionist identification of being in an offside position; and yet, as the laws clearly state, being in an offside position does not necessarily cause an offside offense. Finally, it can be observed that because of this noncorrelation of offside position and offside offense, and even with VAR, judgments of offense are still be prone to an element of human subjectivity and interpretation: clearly different people have different game spirits.

If complexity is the root of the problem, then what of the second question: how could the laws be alternatively formulated to reduce complexity? In answering, it is tempting to propose instead the *minimal* possible definition. So, as a first go, how about: *an offside offense is caused by a player gaining an advantage from being in an offside position*. No need to mention "becoming involved in active play," who touched the ball and when, deliberately or accidentally, interfering with an opponent, having to define a "save" or deal with "line of sight" issues: the referee simply has to decide if, in her opinion, a player was gaining an (implicitly unfair) advantage over an opponent by being in an offside position.

This simplified definition could clearly resolve the situation described above in favor of the defender; but it would also allow the referee to use her skill and judgment to decide all those marginal situations where the attacking player has a toenail in an offside position and, after an interminable delay, a goal is ruled out by a narrative-denying VAR. One problem with video recording and image processing technology is that it is so much more sensitive than the human eye can detect in real time: decisions are being affected by the thickness of the lines being drawn on the display, or which frame of the video was being used (the laws even had to asterisk "the moment the ball is played or touched" as "the

first point of contact with the ball," since in the duration that a player's foot was in contact with the ball while kicking it, the teammate could move from a micro-onside position to a micro-offside one).

While this might go against the "spirit" of the fundamentalist interpretation of the rules adherence brigade, for whom an offside position *is* an offside offense, no matter what the laws actually say, it is consistent with the "spirit" of those for whom such literal hair-splitting is not providing any sort of meaningful advantage and makes no material difference to the actual outcome of the players' actions. Recall this is a socially constructed fact based on an arbitrary line defined by a convenient point of reference—compare the definition of equator separating north and south (a line based on physical reality) with the definition of the prime meridian (0° longitude) separating east from west (a line based on conventional agreement).

Therefore, the "spirit" of the "legal" offside definition begs a question: how far can a player be in an offside position before she is causing an offside offense by gaining an advantage. And this question is one of those questions encountered in philosophy classes known as a Sorites Paradox [4]: the problem of vague predicates being coerced into a binary classification by applying an arbitrary metric.

The classic example of a Sorites Paradox is to suppose that there is a pile of sand, which an observer agrees is a pile of sand. The philosopher removes one grain from the pile, and asks again: is this a pile of sand? The observer affirms, and the philosopher removes another grain; and the process is repeated until the question becomes: how many grains must be removed before there is not a pile of sand, or is one grain of sand also a pile? The general problem is caused by small changes which are individually indistinguishable to the human eye and mind, while the accumulation of these small changes is distinguishable and brings about a "flip" in the binary—from state X to *not-X* [5].

It has been argued that a Sorites Paradox involves a question which is, in some part or sense, "unknowable" (epistemologically [6]), yet at the same time, the answer is, or can be, "known" (practically). One way to determine whether a player is "gaining an advantage by being in an offside position" is by appealing to consensus, so that the truth of a vague predicate is assigned by group usage. The use of consensus suggests that "gaining an advantage by

being in an offside position" can be determined by the proportion of a group that believe this to be the case. There are various procedures for doing this for different types of questions in different circumstances [7]–[10], and in this context it might be objected that the answer is overly dependent on the constitution of the group. In a partisan setting involving a competitive football match, the supporters of the defending team might almost certainly claim an offside offense, while the supporters of the attacking team might prefer to give the benefit of the doubt (at least until it is their team that is defending).

So, while there is no doubt or dispute over the truth of the situation, the possibility of community bias can produce divergence over the meaning [11]: one group's offside is another group's "not in the spirit of the game." Moreover, this meaning is both empirical and contextual: it is based on observation more than application of the arcane offside law and can be biased by numerous contextual factors (the teams, the players, preceding events, etc.). It might, therefore, be concluded that these situations are essentially indeterminate: it cannot be definitively asserted that a player is, or is not "gaining an advantage" by being in this offside position.

But what if we insisted that we map this vague predicate into a crisp binary proposition. We might then observe that within a group, individual members would disagree over whether any one particular situation can be defined as an "gaining an advantage by being in an offside position," but—supposing the group members could be put in a position of complete neutrality—we might also observe that there is a reasonably well-defined boundary, which in itself is "unknowable" either individually or collectively [12], but where a significant majority of the group "just know" that an offside position one side of the boundary is gaining an unfair advantage and should be penalized as an offside offense, and the other side is not.

This provides a possible answer to our third question: what procedure could be used to make a decision under this formulation of the offside rule, assuming that we still wanted to keep the referee/ VAR in control and do not want to crowdsource the decision in "real time." One approach would be to use crowdsourcing, but take decisions offline: this would then be data. It might then be possible to use this data to train a machine learning system to discriminate between offside positions in which the attacking player is gaining an unfair advantage and those where she is not. And this leads to the possibility of the absolutely minimal offside law (see [3]): "an offside position is an offside offense if the computer says it is." This is, perhaps, the ultimate reductionist form of procedural justice [13].

However, this minimal definition leads directly to a consideration of the final question raised earlier: if it were possible to formulate laws involving vague predicates and adjudicate them this way, what would be the implications of such minimalist formulations for soft laws and even for "hard" laws? The possible implications are threefold: 1) does possibility imply desirability; 2) does possibility imply infallibility; and 3) does possibility imply accountability?

The answer advanced here, to all three questions, is "no." On the first implication (possibility implies desirability), while it might be acceptable in a sporting arena, since the situation described at the start of this article affects nothing more serious than local Pyreneean bragging rights, it is not necessarily the case that this should be extended to other domains of human enterprise. Take, for example, the Republican-controlled Senate blocking Barack Obama's appointment of Merrick Garland to the U.S. Supreme Court on specious grounds, while ushering through a deeply conservative appointment to replace liberal judge Ruth Bader Ginsburg. Even the harshest critics of such partisan hypocrisy accept that no rules were broken, but that they were interpreted and applied inconsistently and not in "the spirit" (those rule spirits again) that the original drafters intended.

That is not to say that interpretation is not valid, but we need that interpretation to be exposed to scrutiny. There is clearly still a role for logical representation and formal inference as well as empirical precedent and the human interpretation of vague predicates in computer-aided legal drafting and application, rather than hiding these processes. Algorithmic justice and algorithmic governance based on the principle "it is because the computer says it is" is surely neither just nor legitimate.

On the second implication (possibility implies infallibility), it would be interesting to issue the following "j'accuse" challenge to data science and artificial intelligence [14]: when the crunch came, and the pandemic effluent hit the societal fan, for all the vaunted promises of big data and machine learning, they failed. Moreover, they failed because so much creativity and innovation of science and technology

was focused on profit-making and not on pro-social benefits, or the general socio-economic challenges which are the lived experience of ordinary citizens, or the specific actions that would have helped protect them individually and collectively. It might, for example, have been more useful to know that the United Kingdom was ill-equipped to deal with a pandemic because ten years of Conservative government austerity had reduced stocks of protective clothing to an all-time low, rather than knowing about a lack of toilet paper based on sentiment analysis of internet search engine queries and social media feeds.

And when it came to social media platforms, not only did they fail, they arguably made matters worse, because even when they were an effective prevention of serious illness through vaccination, the dissemination of disinformation and misinformation created resistance in a substantial percentage of the population, who rejected that scientific narrative while swallowing (quite literally, in some cases) others (e.g., hydroxychloroquine and ivermectin). And they failed because they too used machine learning to raise addiction and indentured servitude in the pursuit of profit above the pro-social benefits of common knowledge, collective action and the public interest [15].

ON THE THIRD IMPLICATION (possibility implies accountability), the answer must again be "no," because if anything this possibility increases non-accountability. Although there are attempts to make the machine justify itself through explainability or answerability [16], the risk is that algorithmic governance of this self-defining sort would allow certain politicians not just to be one step further from the truth [17], but also to be a giant leap further removed from taking responsibility for their actions; especially when those actions are demonstrably not in the public interest, unlawful, suppress human rights, break international law or renege on international treaties, and are accompanied by an autocratic tendency which actively seeks to avoid scrutiny and accountability by the press, the courts, judicial review, or public inquiry. For example, the U.K. Government, in particular, has made a great deal out of "following the science" during the pandemic [18]—which is willfully misleading: science advises, politics decides; scientists can take responsibility for poor advice, but politicians must take responsibility for wrong and harmful decisions. How much easier their

defense would be, if they could claim that they were just "following the algorithm," when they should be arraigned for reckless endangerment, misuse of public funds, and even crimes against humanity. ∎

■ References

[1] IFAB/FIFA. *Laws of the Game 20/21*. Accessed: Nov. 16, 2021. [Online]. Available: https://digitalhub.fifa.com/m/5371a6dcc42fbb44/original/d6g1medsi8jrrd3e4imppdf.pdf

[2] P. Suber, *The Paradox of Self-Amendment*. Bern, Switzerland: Peter Lang, 1990.

[3] R. Singh, A. Bodhe, P. Kanuparthi, A. Ananthakrishnan, and J. Pitt, "From classification to definition: The changing nature of human adjudication," *IEEE Technol. Soc. Mag.*, vol. 38, no. 4, pp. 55–62, Dec. 2019.

[4] C. Barker, "Vagueness," in *Concise Encyclopedia of Semantics*, K. Allan, Ed. Amsterdam, The Netherlands: Elsevier, 2009.

[5] R. Collins, "On the borders of vagueness and the vagueness of borders," *Vassar College J. Philosophy*, vol. 5, no. 5, pp. 30–44, 2018.

[6] T. Williamson, "Inexact knowledge," *Mind*, vol. 101, no. 402, pp. 218–242, 1992.

[7] J. Surowiecki, *The Wisdom of Crowds*. Boston, MA, USA: Little, Brown, 2004.

[8] J. Ober, "Democracy's wisdom: An Aristotelian middle way for collective judgement," *Amer. Political Sci. Rev.*, vol. 107, no. 1, pp. 104–122, 2013.

[9] J. Pitt, "Interactional justice and self-governance of open self-organising systems," in *Proc. IEEE SASO*, Oct. 2017, pp. 31–40.

[10] J. Pitt, A. Nowak, T. Michalak, W. Borkowski, and R. Vallacher, "Knowing what the bits know: Social influence as the source of collective knowledge," in *Proc. 2nd Int. Workshop Agent-Based Modelling Hum. Behav.*, 2020, pp. 1–7. [Online]. Available: http://abmhub.cs.ucl.ac.uk/2020/

[11] H. Arendt, *The Life of the Mind*. Boston, MA, USA: Houghton Mifflin Harcourt, 1978.

[12] T. Williamson, *Vagueness*. London, U.K.: Routledge, 1994.

[13] J. Rawls, *A Theory of Justice*. Boston, MA, USA: Harvard Univ. Press, 1971.

[14] T. Snyder, *Our Malady*. New York, NY, USA: Random House, 2020.

[15] J. Pitt, K. Michael, and R. Abbas, "Public interest technology, citizen assemblies, and performative governance," *IEEE Technol. Soc. Mag.*, vol. 40, no. 3, pp. 6–9, Sep. 2021.

[16] A. Adadi, and M. Berrada, "Peeking inside the black-box: A survey on explainable artificial intelligence (XAI)," *IEEE Access*, vol. 6, pp. 52138–52160, 2018.

[17] P. Oborne, *The Assault on Truth: Boris Johnson, Donald Trump and the Emergence of a New Moral Barbarism*. London, U.K.: Simon & Schuster, 2021.

[18] J. Calvert and G. Arbuthnott, *Failures of State: The Inside Story of Britain's Battle With Coronavirus*. London, U.K.: Mudlark, 2021.

Jeremy Pitt is a Professor of Intelligent & Self-Organising Systems with the Department of Electrical & Electronic Engineering, Imperial College London, London, U.K. He is a Fellow of the British Computer Society (BCS), the Institute for Engineering and Technology (IET), and a member of IEEE. He is currently the Editor-in-Chief of *IEEE Technology and Society Magazine*.

■ Direct questions and comments about this article to Jeremy Pitt, Department of Electrical and Electronic Engineering, Imperial College London, London SW7 2BT, U.K.; j.pitt@imperial.ac.uk.