

MAXIM TOKENS AND MAXIM TYPES: WHY UNIVERSALIZABILITY NEED NOT BE UNIVERSAL

SAMUEL KAHN

Abstract. In this article, I argue that Kant's Categorical Imperative applies to maxim tokens rather than to maxim types. The article has three main parts. In the first, I explain my thesis. In the second, I argue for it. In the third, I argue, further, that, if my thesis is correct, then tokens of different maxim types can have different deontic statuses for different agents.

Keywords: Kant's ethics; Kantian ethics; Categorical Imperative; maxim; maxim type; maxim token; act token; act type; universalizability.

The article has three main parts. In the first, I explain my thesis. In the second, I argue for it. In the third, I argue, further, that, if my thesis is correct, then tokens of different maxim types can have different deontic statuses for different agents.

MAXIM TOKENS AND MAXIM TYPES

Consider the lying-promise maxim (LPM), "I will tell a lying promise in order to get some ready money", from Kant's famous example in the *Groundwork to a Metaphysics of Morals*. There are three related things that I want to draw attention to about the LPM. These three things, in turn, will help to explicate CI-TOKEN (i.e., the thesis that Kant's Categorical Imperative applies to maxim tokens rather than to maxim types).

First, the LPM can be adopted on many different occasions by the same agent. For example, suppose that Kat, who is rather unscrupulous, is facing eviction because she is behind on her rent. Kat decides to defraud one of her friends, and, accordingly, she adopts the LPM. She receives the money; pays her rent; and loses her friend. Shortly thereafter, Kat joins a religious cult. She gets

Samuel Kahn ✉

ever more involved with the group. Ultimately, she undergoes a conversion, living on a commune with this group in relative comfort. Time passes, and, over the years, Kat's convictions change, along with her principles, until she reaches a point at which she not merely would not tell a lying promise when in financial exigency, but she would not do so on moral, rather than, say, prudential, grounds. The point is that Kat no longer wills the LPM and, indeed, if someone were to ask her about doing so, Kat strenuously would condemn such a principle (and would condemn her former self for having acted on it). After a decade of living in this way, however, disaster strikes. The guru leading Kat's cult dies. Various details emerge, suggesting that this guru was not the morally perfect saint that she projected herself to be. Devastated, Kat leaves the group, falling into general malaise. More years go by. Kat does not revert back to her former, unscrupulous ways. But, she no longer has the moral strength she had during the height of her time on the commune. Finding herself in financial exigency again, she decides, once more, to tell a lying promise--and, accordingly, Kat adopts the LPM. The purpose of this vignette is not to condemn Kat or to praise her; it is merely to make plausible the claim that the LPM can be adopted on many different occasions by the same agent (in this case: by Kat, before and after her involvement with the cult).¹

Second, the LPM can be adopted by different agents. To make this plausible, we can imagine another agent, Mat, halfway across the world from Kat, who finds himself in danger of having his car repossessed because he has missed several payments. Although Mat is geographically distant from Kat, his situation is similar: both are in financial exigency, and, we may suppose, Mat decides that the best solution to his problem is to make a lying promise. Thus, Mat, like Kat, adopts the LPM. If this is possible, then, as asserted in the first sentence of this paragraph, different agents can adopt the LPM.²

Third, the LPM is one of many different kinds of lying-promise maxim. For example, Nat, who lives on the other side of Mat from Kat, might adopt a maxim to tell a lying promise in order to get a smartphone or a laptop rather than money ("I will tell a lying promise in order to get Mat to let me "borrow" his laptop"). Nat,

¹ Some complications arise here on account of the fact that it seems at least *prima facie* plausible that an agent can act on a maxim on more than one occasion, without rejecting the maxim and then adopting it again. For example, Kat might tell multiple lying promises during her unscrupulous phase, before she undergoes her conversion and rejects the LPM. We might wonder, then, whether maxim tokens are individuated by action tokens (in which case Kat's multiple lying promises during her unscrupulous phase count as multiple maxim tokens), or whether maxim tokens are individuated by adoption tokens (in which case these multiple lying promises involve a single maxim token). For present purposes, this question does not need to be settled, so I propose to ignore it (indeed, it was precisely to avoid grappling with this question that I added so much detail to the vignette about Kat: this vignette involves two tokens of the LPM on any plausible account of maxim token individuation).

² It is perhaps worth pausing to point out that, if it is not possible for Mat to adopt the LPM -- if, for some reason, only Kat can adopt the LPM -- then a large chunk of work on Kant's ethics, including Kant's own, will have to be thrown out.

who lives on the other side of Kat from Mat, might adopt a maxim to tell a lying promise in order to get someone to marry her (perhaps Jat wants a Green Card), or in order to get a promotion at work (perhaps Jat desperately wants more social capital). Thus, the LPM is one of a family of maxims that can be grouped by virtue of the fact that they all involve telling lying promises for one reason or another.

We can use the first two points to make sense of the distinction between a maxim token and a maxim type. A maxim token is a particular instance of the adoption of a maxim type. Thus, in adopting the LPM on two separate occasions, Kat adopts two tokens of the type. If we add Mat's adoption into the mix, then we have three tokens of the type.

We can use the third point in order to clarify the nature of CI-TOKEN and, in particular, the idea that the Categorical Imperative applies to maxim tokens rather than anything else. This idea amounts to saying that we must adopt a bottom-up, as opposed to a top-down, approach to the derivation of duties. Let me explain.

A top-down approach to the derivation of duties will begin either with claims about the deontic status of general kinds of maxims, which subsume many different maxim types, or with the deontic status of maxim types themselves, and then infer the deontic status of maxim tokens.

For example, someone might use the Categorical Imperative to argue that it is always impermissible to adopt lying-promise maxims, no matter what the end of the lying-promise happens to be, whether money (Kat and Mat), a smartphone (Nat), or marriage or a promotion (Jat).³ From this, it then can be inferred that, any

³ Technically, in Kant's and Kantian ethics, maxims do not have deontic properties: actions have deontic properties based on the (non-normative) properties of the maxims on which they are performed. For example, according to the universalization formulations of the Categorical imperative, act X is impermissible if, but only if, the corresponding maxim is not universalizable. Similarly, according to the humanity formulation of the Categorical Imperative, act X is impermissible if, but only if, the corresponding maxim fails to use humanity in any person at the same time as an end. The shorthand, of ascribing deontic properties to maxims directly, is based on the following biconditional: maxim M has deontic property D if, but only if, any action performed on the basis of M has deontic property D. The deontic properties of the action are then determined using biconditionals like the ones in previous paragraph of this note. Telescoping these various biconditionals, we get the conventional (even if, technically, incorrect) biconditionals about the deontic properties of maxims that lie in the background of most recent work in Kantian ethics, such as: maxim M is impermissible if, but only if, it is not universalizable. Questions arise here, and refinement of the biconditional in the final sentence of the previous paragraph might be needed, in the case of acts performed on the basis of more than one maxim. For example, if it is possible for an act to be performed on the basis of several maxims, and if these several maxims have different properties, then it is unclear what the deontic status of the act will be, and this feeds back into unclarity about the biconditionals in the previous paragraph. Perhaps a refinement like the following will work: maxim M has deontic property D if, but only if, any action performed solely on the basis of M has deontic property D. However, it is unclear. These complications can be sidestepped if an act cannot be performed on the basis of more than one maxim. That is, it might be argued that act tokens are individuated on the basis of the maxim on which they are performed, and it might be argued that, when more than one maxim seems to be at play, really, this should be construed as a case of a conjunctive maxim (or something along these lines). If this is correct, then the biconditionals in paragraph 2 of this note can be affirmed without refinement and

time an agent acts on such a maxim, she is acting impermissibly. Thus, if someone adopts this strategy, then, from her conclusion that lying-promise maxims in general are impermissible, she will infer, in one fell swoop, that Jat's, Kat's, Mat's, and Nat's different lying-promise token maxims are all impermissible.

Alternatively, someone who adopts a top-down approach – or, perhaps more accurately, a middle-down approach – might, instead, use the Categorical Imperative to argue that the LPM, the maxim type, is impermissible, regardless of whether lying-promise maxims in general, like Nat's and Jat's, are so. From this it then can be inferred that any tokens of the type (Kat's and Mat's several maxim tokens) are impermissible. Thus, if someone adopts this strategy, then, from her conclusion that the LPM (the maxim type) is impermissible, she would infer that Kat's and Mat's tokens of the LPM are all impermissible. But, on this view, an additional argument is needed to show that Jat's and Nat's maxim tokens are impermissible.

A bottom-up approach, by way of contrast, reverses the direction of all the inferences described in the previous two paragraphs. A bottom-up approach begins with the individual maxim tokens and infers, from the properties of these tokens, to conclusions about the corresponding types and kinds. Thus, a bottom-up approach begins with five separate arguments to show that Jat's, Kat's, Mat's, and Nat's respective maxim tokens are all impermissible. (Of note: precisely because Kat adopts two maxim tokens, two distinct arguments are needed, on this approach, to show that each one is impermissible, this notwithstanding the fact that they are tokens of the same type.) To get to the conclusion that the maxim type is impermissible, or that, in general, lying-promise maxims are impermissible (regardless of the end), some sort of generalization from tokens (bottom) to type (middle) or kind (top) is necessary.⁴

without worrying about acts performed on the basis of more than one maxim (because there are no such acts). Two caveats are worth mentioning here. First, this last position (namely: that an act cannot be performed on the basis of more than one maxim) is consistent with there being a maxim “stack” (of the kind Onora O’Neill, among others, discusses) behind any given act. In such cases, the permissibility of the adoption of any given maxim in the stack is determined by the maxim on the basis of which it was adopted, and the permissibility of the act is determined on the basis of the lowest maxim in the stack. (This has the *prima facie* odd, but *ultima facie* intuitive, result that a universalizable maxim can be adopted on the basis of a non-universalizable maxim and, therefore, a permissible maxim can be adopted in an impermissible way.) Second, this last position is also consistent with the idea that multiple act tokens can be performed on the basis of a single maxim token. Just because act tokens are individuated on the basis of the maxim tokens on which they are performed does not entail that maxim tokens are individuated on the basis of the act tokens that they are instantiated in (see note 1 above). Fortunately, these puzzles can be set aside for the moment: the various positions we might take on them are neither entailed by, nor entail, CI-TOKEN.

⁴ Alternatively, using generalities, we can establish a deliberative presumption that tokens of a type are impermissible (one can even do this by showing that tokens of a type that belongs to a given kind are, in general, impermissible: e.g., by showing that maxims that involve lying-promises, regardless of the end, are, in general, impermissible). But, the point is that neither of these argument strategies (i.e., neither the argument strategy described in the sentence to which this note is appended, nor the argument strategy described in the first sentence of this note) allows for a direct, deductive

The reason this is relevant is that, if CI-TOKEN is correct--if the Categorical Imperative applies to maxim tokens rather than to maxim types--then a bottom-up approach to the derivation of duties must be adopted and, correlatively, top-down (or middle-down) approaches must be rejected, this despite the fact that the latter are the standard in discussions of Kant's and Kantian ethics.⁵

MAXIM TOKENS, NOT MAXIM TYPES

The argument I am going to make for CI-TOKEN, which is an adumbrated version of an argument I have made elsewhere, is based on the ought implies can principle (OIC), a meta-ethical principle that bridges the moral-modal gap.⁶

Kant's commitment to OIC is well-known. In a famous example from the *Critique of Practical Reason*, Kant suggests that, even if someone were faced with death, he would judge it possible to overcome his love of life in order to act morally. Kant infers from this example that, in general, a person "judges that he can [do] something on the basis of the fact that he is conscious that he ought [to do] it."⁷ The basic point is that, according to Kant, if an agent has a duty to D, then she is able to D, or, in its contrapositive form, if an agent is unable to D, then she does not have a duty to D.

There is considerable debate about how to understand OIC, and these debates are both philosophical and exegetical. Philosophically, there are questions about the meaning of 'ought', the meaning of 'implies', and the meaning of 'can'. For example, when it comes to the meaning of 'implies', some maintain that OIC is about logical entailment; others maintain that it is about presupposition; and yet others maintain that it is about conversational implicature. These three interpretations come apart because they allow for different inferences in cases of inability (when the contrapositive form of OIC is at play). Although some reject OIC on any interpretation, others maintain, on the basis of these kinds of differences (e.g., differences about what may be inferred in cases of inability), that OIC is defensible

inference that a given token of the type is impermissible: these arguments establish only rebuttable presumptions to that effect. With a top-down approach, by way of contrast, we may appeal to an inference rule like universal-instantiation in order to infer the properties of the token from the type.

⁵ To be clear: the way we can know that bottom-up approaches are not the standard in discussions of Kant's and Kantian ethics is precisely what was pointed out in note 4, namely: with a bottom-up approach, we get only rebuttable presumptions about maxim tokens from arguments about maxim types, whereas with top-down (or middle-down) approaches, deductive inference rules like universal-instantiation may be used to infer the deontic status of a maxim token from the deontic status of a maxim type--and what is standard in discussions of Kant's and Kantian ethics is to treat such inferences as deductively valid.

⁶ For an extended version of this argument, see section 4 of Samuel Kahn, "Individual Maxim Tokens, not Abstract Maxim Types", in *Kantian Review*, 2024 online, pp. 1–17.

⁷ KpV, AA 05: 030.33–34, my translation.

on some, even if not on other, interpretations.⁸ Exegetically, there are then questions about which version of OIC Kant is committed to.⁹

For present purposes, it suffices to compare two different versions of OIC. According to OIC-token, a duty to D implies the ability to perform the act token, D. According to OIC-type, a duty to D implies the general ability to perform actions of type D, even if not the token D. To see how OIC-token and OIC-type come apart, suppose that, out of the corner of her eye, Leona sees a small child, Aitana, struggling to stay afloat in a pond. Suppose, further, that, although, in general, Leona has the ability to run over to the pond to save Aitana, in this particular instance, Leona's ability is masked (perhaps Leona gave blood earlier in the day, and, if she exerts herself now, she will faint midway to the pond).¹⁰ In other words, Leona can perform the act type of running over to the pond to save a child, but she cannot perform this particular token of the type--she cannot save Aitana. In this case, OIC-token nullifies the duty of rescue that Leona has to Aitana, whereas OIC-type does not.¹¹ The question, now, is: which version of OIC, OIC-token or OIC-type, is Kant committed to?

Kant answers this question in the *Religion within the Boundaries of Mere Reason*:

[D]uty charges an agent unconditionally: he ought to remain true to himself; and from this he rightfully infers: he must also be able to do so, and his will is thus free...freedom, according to which the action as well as its opposite must be, in the instant of occurrence, within the power of the subject.¹²

Kant says that the action which a subject is obliged to perform must be, by virtue of that obligation, within the power of the subject *in the instant of occurrence*. From this it may be inferred that Kant is committed to OIC-token: if an agent ought to D, then she is able to perform the act token D (not merely types of that token).

⁸ I canvass these debates in part 1 of Samuel Kahn, *Kant, Ought Implies Can, the Principle of Alternate Possibilities, and Happiness*, Maryland, Lexington Press, 2018.

⁹ I canvass these debates in chapter 2 of Samuel Kahn, *Kant, Ought Implies Can, the Principle of Alternate Possibilities, and Happiness*, Maryland, Lexington Press, 2018; in sections 1 and 2 of Samuel Kahn, "Some Contemporary Issues about Ought Implies Can: Where Does Kant Fit in?", in *Jahrbuch für Recht und Ethik*, Vol. 31, Issue 1, 2023, pp. 187–207; and in section 4 of Kahn, Samuel "Individual Maxim Tokens, not Abstract Maxim Types", in *Kantian Review*, 2024 online, pp. 1–17.

¹⁰ This example is (very) loosely based on a similar example from Markus Kohl, "Kant and 'Ought Implies Can'", in *The Philosophical Quarterly*, Vol. 65, No. 261, 2015, pp. 690–710.

¹¹ Some care is needed here, modulo the issues raised above (in the previous paragraph, the one to which note 9 is appended), regarding the interpretation of OIC. For example, if 'implies' in OIC is understood as conversational implicature, then, in the scenario envisioned in the paragraph to which this note is appended, there is (merely) an over-ridable conversational norm against the assertion of of Leona's duty of rescue. These issues will be overlooked here: the most important issue for present purposes is the type/token distinction.

¹² RGV, AA 06: 049.30–50.27, my translation.

The reason this is relevant for present purposes is that Kant's commitment to OIC-token militates in favor of CI-TOKEN, the thesis that the Categorical Imperative applies to maxim tokens rather than to maxim types. To see why, suppose that CI-TOKEN is false; suppose that the Categorical Imperative applies to maxim types. Then, following the top-down approach, the deontic status of a maxim token is inferred from the deontic status of the corresponding maxim type. Thus, to return to the example from the previous section, the impermissibility of Kat's and Mat's token adoptions of the LPM is inferred from the impermissibility of the corresponding maxim type. From this it may be seen that, if CI-TOKEN is false, then abilities at the instant of occurrence cannot have any bearing on obligations, contradicting OIC-token. It may be concluded that Kant's commitment to OIC-token supports CI-TOKEN.

It might be thought that CI-TOKEN can be pried apart from considerations about OIC-token. For example, it might be thought that, by appeal to excusing conditions, one could affirm OIC-token while rejecting CI-TOKEN. That is, some might think that the Categorical Imperative applies to maxim types and that OIC-token is nonetheless true, because OIC excuses an agent from her obligations, while leaving those obligations in place for other agents.

However, this thought is based on a mistake: this thought mistakes the nature of an excuse. Excuses, if such there be, leave obligations intact--excuses nullify blame, not obligations.¹³ The idea behind an excuse is that, if an agent has an excuse, then she has failed to fulfill an obligation, but she is not to blame for this failure. From this it may be seen that the thought in the previous paragraph subtly shifts us away from OIC and toward a related principle, the principle that blame implies can (BIC).

There are those who argue for BIC over OIC: they maintain that OIC should be rejected for one reason or another, and they maintain that the intuitions that support OIC can be explained by BIC, which is seen as a weaker principle.

However, that strategy will not work exegetically: Kant's commitment to OIC is not in question, nor, as evidenced even by the short quotation from the *Religion within the Boundaries of Mere Reason* reproduced above, should it be. So, any strategy that requires us to jettison OIC in favor of BIC will not work from within the framework of Kant's ethics: such a strategy cannot be reconciled with the texts. Moreover, although a full defense of OIC cannot be undertaken here, it should be added that the retreat from OIC to BIC is philosophically problematic. To see why, note that, when we consider, for example, the suffering of people who died under intolerable conditions in years gone by (e.g., Jews subject to medical experimentation in the Holocaust)--suffering that, by virtue of being in the past, we are unable to alleviate in any way--it is not merely that we are not to blame for not relieving their suffering, it is that we do not have a duty to relieve their suffering.

¹³ Paul M. Hughes and Brandon Warmke, "Forgiveness", in Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Summer 2017 Edition), sections 2.1–2.2.

The best explanation of this is based on two facts: (1) we are unable to relieve the suffering in question, and (2) this inability nullifies, rather than merely excuses us from, any corresponding duty, as per OIC. To bolster (2), suppose that it is possible for me to snap my fingers and thereby to help these long-dead people. In that case, it seems plausible that, all else being equal, I have a duty to snap my fingers and help them, and, absent a plausible excuse, I am to blame if I fail to do so. If this is correct, then OIC is the best explanation of the fact that, given (1), I do not have any such duty.¹⁴

To summarize: not only Kant's, but any defensible Kantian, ethics is committed to OIC, not BIC; the specific version of OIC in question, if we are contrasting OIC-type and OIC-token, is OIC-token; and commitment to OIC-token supports CI-TOKEN, the thesis that the Categorical Imperative applies to maxim tokens, not maxim types. Thus, Kant's, and any defensible Kantian, ethics should be committed to CI-TOKEN, which is what was to be shown.¹⁵

MAXIM TOKENS AND AGENTS' WILLS

Having explained CI-TOKEN (in section 1) and argued for it (in section 2), the goal of this final section is to show that, if CI-TOKEN is correct, then different tokens of the same maxim type can have different deontic statuses. That is, although, for the most part, when applying the Categorical Imperative to maxim tokens, tokens of the same type will have the same deontic status for most agents, it is logically, physically, and really possible that there will be instances in which this is not the case.

To see this, we will proceed in two stages. First, the three best-known formulations of the Categorical Imperative will be introduced. Second, I will make two observations about how these different formulations of the Categorical Imperative, when applied to maxim tokens, can yield different, agent-relative results, this despite being objective (i.e., universal, non-agent-relative) moral standards. The observations will be about: abilities, practical judgment, beliefs, and, specifically, how these things inform the Categorical Imperative.

¹⁴ I owe this argument strategy to Peter Graham, "'Ought' and Ability", in *The Philosophical Review*, Vol. 120, No. 3, 2011, pp. 337–382.

¹⁵ This is not the only argument that can be given in favor of my thesis. For example, one also might appeal to Kantian ideas about autonomy. Kant famously maintains that his focus on autonomy is what distinguishes his ethics from all other systematic ethics, and he proposes to call the Categorical Imperative the principle of the autonomy of the will on precisely these grounds (GMS, AA 04: 432.25–433.11 and 440.16–32). However, if an agent's own will is going to be lawgiving for her, then it must be that agent's own token willings that are willed as universal laws (or laws of nature), or that are supposed to serve as legislation in a possible kingdom of ends. Thus, considerations of autonomy, a central notion in Kant's and any plausibly Kantian ethics, also support my thesis, that the Categorical Imperative should be applied to maxim tokens, not maxim types. For further discussion, see my Samuel Kahn, "The Apple of Kant's Ethics: i-Maxims as the Locus of Assessment", in *Pacific Philosophical Quarterly*, Vol. 4, No. 3, 2022, pp. 559–577.

The most frequently used formulations of the Categorical Imperative in derivations of duties, both in Kant's and in Kantian ethics, are the Formula of Universal Law, the Formula of a Law of Nature, and the Formula of Humanity:

1. Act only according to that maxim through which you at the same time can will that it become a universal law.¹⁶
2. Act thusly, as if the maxim of your action should become, through your will, a universal law of nature.¹⁷
3. Act thusly, that you use humanity, just as much in your person as in the person of any other, always at the same time as an end, never merely as a means.¹⁸

Without taking a stand on the details of how to interpret these different formulations of the Categorical Imperative, note that, if CI-TOKEN is correct and, thus, these formulations are applied to maxim tokens (rather than maxim types), then the first enjoins each individual agent to act only on maxims that she, in particular, can will at the same time as a universal law; the second enjoins each individual agent to act only on maxims that can become universal laws of nature through her particular will; and the third enjoins each individual agent to act so that she, in particular, uses humanity always at the same time as an end, never merely as a means. The point that will be made now is that, precisely because, if CI-TOKEN is correct, these formulations are particularized in this way, the implications of these formulations are so as well and, therefore, the properties of maxims and actions will be relativized to a particular agent. From this it may be concluded that tokens of different maxim types might have different deontic statuses for different agents. This point will be made, again, by means of two observations (observations about: abilities, practical judgment, and beliefs).

The first observation returns us to OIC-token. Different agents have different im/permissions on account of their different abilities. For example, an agent who has the ability to snap her fingers and end world hunger has a duty to do so; an agent who has the ability to snap her fingers and kill everyone within a 15-mile radius has a duty *not* to do so; and agents like the rest of us, who have no such special abilities, are generally permitted to snap our fingers or not, as we will. This illustrates how, if CI-TOKEN is correct, im/permits will be relativized to agents.

It might be objected that the examples used in the previous paragraph are unrealistic: nobody has the ability to end world hunger by snapping their fingers, and nobody has the ability to kill everyone within a 15-mile radius by snapping their fingers. Some might try to use this to push against the idea that im/permits are relativized to agents: they might maintain that all of us have, more or less, the same abilities, and, therefore, any variation is likely to be quite ephemeral.

¹⁶ GMS, AA 04: 421.06-08, emphases omitted, my translation.

¹⁷ GMS, AA 04: 421.18-20, emphases omitted, my translation.

¹⁸ GMS, AA 04: 429.10-12, emphases omitted, my translation.

There is something to this objection: realistically, our abilities fall within a fairly constrained range and, for this reason, OIC-token does not ground grand differences in im/permits of the kind that might arise when thinking about, e.g., the various characters in Greek mythology (i.e., the abilities of mere mortals compared to those of demigods and Gods). But, there is variation in our abilities all the same, and this variation then manifests in differences in our im/permits, differences which, so long as we are focused on the here-and-now, will, at least sometimes, be quite marked. For example: differences in levels of wealth ground differences in regard to duties of beneficence, and similar things can be said about differences in power, influence, and expertise. So, a direct appeal to OIC-token, conjoined with empirical knowledge about differences in our abilities, entails that tokens of different maxim types have different deontic statuses for different agents, which is what was to be shown.

The second observation takes us into casuistry and the place for practical judgment in Kant's and Kantian ethics. As Rawls points out, there are various facts, which he calls burdens of judgment, that can lead to reasonable disagreement among reasonable people. These facts include: (i) the evidence bearing on a case can be complex and difficult to evaluate; (ii) there can be reasonable disagreement about how different kinds of evidence should be weighted; and (iii) our concepts are vague and subject to difficult cases.¹⁹ Kant also was aware of this, and he mentions it in various places, especially in the *Metaphysics of Morals*. For example, consider the following remarks, from a discussion of the demandingness of the duty of benevolence:

I should sacrifice a part of my wellbeing to others, without hope of return, because it is a duty. But it is impossible to give determinate boundaries [to this sacrifice]: how far this [sacrifice] could go. It comes very much to this: what, for everyone, according to his sensibilities, will be his true needs, the determination of which must be left to everyone for themselves. For to promote the happiness of others with the sacrifice of one's own (one's true needs) would be a maxim that would conflict with itself when one makes it to a universal law. From this it follows that this duty [of benevolence] is only a wide one; it has a playroom in it do to more or less, without letting determinate limits to what should be done be given out. -- The law is valid only for maxims, not for determinate actions.²⁰

¹⁹ John Rawls, *Justice as Fairness*, Cambridge, Belknap Press, 2001, pp. 35–36.

²⁰ MS, AA 06: 393.24–35. It is worth comparing these remarks about benevolence with Kant's remarks about suicide at RGV, AA 06: 081.22–25, and also with his remarks about avarice, and about lying and truthfulness, at MS, AA 06: 432–434. Kant seems to allow space for judgment, and for reasonable disagreement, with regard to all of these duties. One thing that is especially notable about these latter passages is that, unlike the one to which this note is appended, they deal with perfect duties.

Kant claims in this passage that a maxim of sacrificing one's own true needs in order to promote others' happiness is not universalizable. But, he also notes that our true needs are vague. Kant might concede that there are clear cases: nobody's true needs include a Rolex or a Maserati. But, his point is that there also are borderline cases: whether I need to buy new shoes, or whether I can put that off for a few months, despite the fact that my present shoes are falling apart, might be less clear. Because our true needs are vague, and because the duty of benevolence does not require us to sacrifice these true needs, the limits of the duty of benevolence are vague. The key is that, according to Kant, the determination of an agent's true needs must be left to the agent herself, and it is a matter of practical judgment. It follows that, on Kant's account, different agents might hold themselves to different standards of self-sacrifice with regard to the duty of benevolence, and there is no room for moral castigation or critique, either from other agents or from agents' own consciences, merely on these grounds.

Some might object that there is something bootstrappy about this. The problem, according to objectors, is that Kant seems to be saying that, when it comes to the limits of the duty of benevolence, an agent's judgment that X is permissible makes it so. For example, if I judge that it is permissible for me to buy a new pair of shoes instead of giving that money to charity, then it is so, and (the objection continues) the reason that it is so is merely that I judge it to be so – this (the objection concludes) is a clear example of hoisting oneself up by one's bootstraps.

However, this objection mistakes what Kant is saying. On Kant's account, the disagreement is with regard to what counts as our true needs. That is, Kant begins with the premise that it is permissible not to promote others' happiness when doing so requires us to sacrifice our true needs. Kant then points out that the limits of our true needs are indeterminate. In other words, what is subject to practical judgment is not whether it is required to sacrifice our true needs for others' happiness (it is not); rather, what is subject to practical judgment is what counts as our true needs. On Kant's account, reasonable people might disagree about this, whence it follows that reasonable people might hold themselves to different standards, and nonetheless equally fulfill the moral law, at least when it comes to what the duty of benevolence requires. Thus, my adoption of a maxim to buy a new pair of shoes might be impermissible (because I judge that these new shoes do not fall within my true needs) whereas another agent's adoption of this maxim might be permissible (because she judges that these new shoes do fall within her true needs)--it can be permissible for one agent to adopt a token of a maxim type even though it is impermissible for another agent to adopt a token of that maxim type--and there is nothing objectionably bootstrappy here. We can make this more perspicuous by considering how it might play out with other duties.

Consider Kant's remarks about suicide in section II of the *Groundwork to a Metaphysics of Morals*. Kant argues that the maxim to kill oneself from self-love

cannot become, through an agent's will, a universal law of nature.²¹ According to Kant, the purpose of self-love is to prolong life, and, therefore, using this self-love to curtail life involves a contradiction. But suppose, now, that we compare two agents, Mariam and Estefy. Mariam, like Kant, believes that self-love is purposive and, more specifically, that its purpose is self-preservation. Estefy, by way of contrast, eschews purposiveness in nature and thinks, further, that self-love is an evolutionary artifact which persists only inasmuch as it promotes the replicability, under normal conditions, of the genes that this self is a vehicle for, whence she concludes that there might be conditions in which a well-functioning vehicle will self-destruct, perhaps even as a result of overwhelming self-love. In this case, Estefy will be able to will a token of Kant's suicide maxim as a universal law of nature, whereas Mariam will not.²²

Alternatively, consider the third example Kant uses in his discussion of the Formula of Humanity: the natural talents example. According to Kant, "there are, in humanity, predispositions to greater perfection," whence it follows that, because the Formula of Humanity tells us to use humanity always as an end, we ought to cultivate these predispositions.²³ But, now we can return to Mariam and Estefy. We may suppose that Mariam, like Kant, believes that we have various predispositions, fixed at conception (if not before), and that a certain subset of these predispositions constitutes our humanity, whence the *Groundwork* duty to promote our natural talents follows (because to use humanity at the same time as an end just is to promote these predispositions). Estefy, however, believes no such thing: she believes that humanity, at least in the Formula of Humanity, is the capacity to set ends, and nothing more. As with the suicide example, Mariam's and Estefy's duties, it seems, will differ. Mariam's beliefs make it impossible for her permissibly to will a maxim of neglect of her natural talents, for she takes her humanity literally to contain various component parts that must be used in various ways, and, thus, such a maxim (i.e., a maxim of neglect) would fail to use her humanity at the same time as an end. Estefy's theoretical framework, by way of contrast, changes the moral landscape: if it is impossible for Estefy to will a maxim of neglect of her natural talents while using her humanity at the same time as an end, it will not be for the same reason.

The two observations made thus far in this section illustrate two different ways in which the results of the Categorical Imperative can differ for different agents if CI-TOKEN is correct: (1) agents with different abilities will have different im/permits because abilities inform what an agent is able to will in conformity with the Categorical Imperative; and (2) agents with different beliefs will have different im/permits because beliefs and practical judgment inform what

²¹ GMS, AA 04: 421.24-422.14.

²² To be more precise: if Estefy is unable to will a token of Kant's suicide maxim as a universal law of nature, it will not be for the same reason as Mariam.

²³ GMS, AA 04: 430.10-17, my translation.

an agent is able to will in conformity with the Categorical Imperative.²⁴ In the previous section of this article, it was shown, in the discussion of OIC-token, that Kant was aware of (1). Some evidence that Kant was aware of (2) was reproduced in this section. Additional evidence may be derived from Kant's highest good argument.

According to Kant, the highest good is possible only if God exists and we are immortal, and he concludes from this that, if we are to fulfill our duty to promote the highest good, then it must be rational to believe in God and immortality, at least for moral purposes. In the *Critique of the Power of Judgment*, Kant appeals to this argument in its contrapositive form, concluding that, if someone is unable to bring herself to believe in God, then she may forgo the duty to promote the highest good:

This proof, which one easily can adjust to the form of logical precision, does not intend to say: it is just as necessary to accept the existence of God as to recognize the validity of the moral law; hence, whoever cannot convince himself of the first, could judge himself to be free from the bindingness of the second. No! Only the intention of the final end, brought to pass through the observance of the latter, in the world (a happiness of rational beings harmoniously corresponding with the observance of the moral law, as the highest and world's best) would have to be given up in that case.²⁵

The most important aspect of this passage, for present purposes, is that Kant's consideration of the connection between beliefs and duties is individualized. That is, Kant does not say that an agent who is unable to bring herself to accept the existence of God must nonetheless promote the highest good because the latter is a duty. Instead, much as we should expect at this point, Kant says that, if an agent is unable to bring herself to accept the existence of God, then *that* agent's duty to promote the highest good is no longer binding--even though this duty remains in place for agents who have no such hang-ups. From this it may be seen that Kant was aware, at some level, that differences in beliefs among individuals have implications for their individual duties.

It might be objected at this point that the view advocated in this final section is at variance with Kant's and Kantian ethics and, further, that this can be seen even from the universalization formulations of the Categorical Imperative reproduced at the beginning of this section. The idea behind this objection is that Kant's ethics and Kantian ethics in general, and the Categorical Imperative in particular, especially in its universalization formulations, are about universal laws that hold for all rational beings, whereas the view advocated in this final section

²⁴ This is not, nor is it intended to be, an exhaustive list of the ways in which the results of the Categorical Imperative can differ for different agents. For example, different psychological proclivities are relevant in this context, as are different longterm (sometimes existential) projects and plans.

²⁵ KU, AA 05: 450.31–451.07.

says that moral laws are individual, the opposite of universal--and if adhering to universal laws that hold for all rational beings entails that CI-TOKEN is false, then so much the worse for CI-TOKEN.

The problem with this objection is that the view advocated in this article, even the thesis that tokens of different maxim types can have different deontic statuses for different agents, does not require giving up on universal laws that hold for all rational beings. On the view advocated here, the Categorical Imperative is taken to be the Supreme Law of Morality, as it applies to imperfectly rational beings (who experience this Supreme Law as an ought)--and it applies, without exception, universally, to all rational beings. The point is that, if we formulate the results of this application process as lists of different im/permits for different agents, then, although these lists will have a great deal of overlap because different agents are very similar, there will be some differences, tracking differences in abilities, practical judgment, beliefs, etc. So, these lists will be individualized, not universal. Thus, we might say: the universalization formulations are universal laws, but what these universal laws tell us to do is to look for universalizable maxims, and universalizability of maxims is not universal, and there is no pressure for it to be so from within Kant's or a plausibly Kantian ethics.