



**University of  
Zurich**<sup>UZH</sup>

**Zurich Open Repository and  
Archive**

University of Zurich  
University Library  
Strickhofstrasse 39  
CH-8057 Zurich  
[www.zora.uzh.ch](http://www.zora.uzh.ch)

---

Year: 2024

---

## **Normative powers without conventions**

Koch, Felix

DOI: <https://doi.org/10.1080/20403313.2023.2297579>

Posted at the Zurich Open Repository and Archive, University of Zurich

ZORA URL: <https://doi.org/10.5167/uzh-259153>

Journal Article

Published Version



The following work is licensed under a Creative Commons: Attribution-NonCommercial-NoDerivatives 4.0 International (CC BY-NC-ND 4.0) License.

Originally published at:

Koch, Felix (2024). Normative powers without conventions. *Jurisprudence*, 15(1):35-47.

DOI: <https://doi.org/10.1080/20403313.2023.2297579>

## Normative powers without conventions

Felix Koch 

Department of Philosophy, University of Zurich, Zurich, Switzerland

### ABSTRACT

What exactly do we need to do in order to make a promise, or to exercise some other normative power? On a view relied on by many philosophers writing on promising, consent, and related phenomena, the answer is that we must communicate a suitable kind of intention. On this view, power-conferring principles assert that specific normative consequences, determined in part by the content of the communicated intention, attach to such communicative acts, and these principles need not be socially practised or accepted to be true. The paper offers a defense of this convention-independent view against the forceful challenge developed by Jed Lewinsohn in ‘The “Natural Unintelligibility” of Normative Powers’. Lewinsohn appeals to action-theoretic considerations to show that the relevant type of communicative act could not be performed under conditions of rationality and full information, and that therefore promissory power and other normative powers require the existence of social rules conferring such powers. The defense of the view targeted by Lewinsohn turns partly on the exact content of plausible constraints on the intelligibility of actions done with a particular aim, and partly on the question just how the social acceptance of power-conferring rules should be thought to matter.

### KEYWORDS

Promising; normative powers; conventionalism; rationality

What exactly does one have to do in order to make a promise? It is tempting to pass over this question fairly quickly when thinking about the morality of promising, and more generally to bypass the question of what amounts to exercising any particular normative power. Jed Lewinsohn’s paper ‘The “Natural Unintelligibility” of Normative Powers’ is a forceful reminder that such complacency would be unwarranted. We are in his debt for developing an inventive and illuminating challenge to what many – in large measure due to Joseph Raz’s influential writings on promising – have regarded as an attractive conception of what the exercise of paradigmatic normative powers consists in, captured by what Lewinsohn dubs the ‘Communication Principle’.

Lewinsohn’s paper raises deep questions at the intersection of normative ethics and the philosophy of action. It is thoroughly rewarding to think through his careful treatment of these questions, and my discussion is bound to be very incomplete. In what

**CONTACT** Felix Koch  [felix.koch@uzh.ch](mailto:felix.koch@uzh.ch)

© 2024 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group  
This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives License (<http://creativecommons.org/licenses/by-nc-nd/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited, and is not altered, transformed, or built upon in any way. The terms on which this article has been published allow the posting of the Accepted Manuscript in a repository by the author(s) or with their consent.

follows, I will briefly restate what I take to be the main goal of his paper and the main premises of his argument (Section 1), and then summarise his argument against the so-called Aim Principle (Section 2), which Lewinsohn goes on to apply against the thesis that promising consists in communicating a certain kind of intention whose content makes reference to the communicative act itself (the Communication Principle). In Section 3, I point to two respects in which we might think of resisting the argument against the Aim Principle, and I articulate one reason to doubt that the argument, even if sound, carries over to the Communication Principle. Finally, Section 4 raises some questions about the positive proposal gestured at in Lewinsohn's paper, to the effect that we owe our normative powers to the existence of suitable social norms. Readers who do not need to be reminded of the shape of Lewinsohn's argument could skip the first two sections.

## 1. The goal and the materials

To all appearances, the kinds of acts that constitute paradigmatic exercises of what philosophers, following Raz, have come to call 'normative powers' – the kinds of acts constituting, for instance, the making of a promise or the giving of a permission – bring about changes in people's deontic situation in ways that makes these changes responsive in a particularly direct fashion to actual or asserted intentions. As Lewinsohn aptly puts it, they are acts that bring about deontic changes 'by fiat'.<sup>1</sup> His paper offers an ingenious argument against the claim that deontic changes can be brought about simply by triggering certain normative principles that, as he sees it, encapsulate the best accounts we have of what it could mean to bring about a normative change by fiat, apart from possible alternative accounts of fiat acts that rely on 'social rules'.

Lewinsohn's complaint is not a normative one. He does not primarily attack the idea that certain fiat acts could have the kind of deontic significance that they are taken to have by those who embrace the relevant accounts of promising, permitting, and so on. Thus he is not concerned with what some philosophers have claimed to find puzzling: how fiat acts, ephemeral as they are, could have deontic significance. His argument cuts deeper. He aims to show that to the extent that we look only to the two fiat-encoding principles he discusses – the Aim Principle and the Communication Principle – it is conceptually impossible to perform the relevant acts, except under conditions of incomplete information or possibly thanks to some failure of rationality. It is in this sense that those acts are 'naturally unintelligible'.<sup>2</sup>

In arguing for this conclusion, Lewinsohn relies on three key premises:

First, a person can perform some act  $\varphi$  with the aim (goal, purpose, intention)<sup>3</sup> of bringing about X (an event or state of affairs) only if she believes that her  $\varphi$ -ing on

<sup>1</sup>As Lewinsohn notes, the term 'normative power' can be, and has been, legitimately employed to express a concept that does not have this implication. In arguing for the unintelligibility of normative powers whose exercise essentially consists in fiat acts, Lewinsohn's paper targets what even those who employ the term in a wider sense would regard as paradigmatic cases of normative powers, or as instances of a paradigmatic (narrow) concept of normative power.

<sup>2</sup>I return in the final section to why the qualifier 'naturally' may be misleading here, to the extent that it suggests that these acts could be intelligible in ways other than by being 'naturally' intelligible. I think that the conclusion supported by Lewinsohn's argument, if successful, is that these acts are unintelligible (under conditions of rationality and adequate information) except to the extent that they have normative consequences in virtue of something other than the normative principles he discusses.

this occasion could in some relevant sense ‘contribute’ to *X*’s coming to pass. This is the Contribution Condition.

Second (roughly stated), where one act-type is a specification<sup>4</sup> of another act-type, a person’s performing an act of the more general type cannot in the relevant sense contribute to her having performed an act of the more specific type. This is the Specification Constraint.

Third, where some fact or set of facts *y* fully and directly accounts for some distinct fact *z*, and where *z* is not overdetermined, any further fact *x* that also at least partly accounts for *z* can do so only by partly or fully accounting for *y*. Call this the Explanatory Constraint.

Relying on these elements, Lewinsohn argues against the two normative principles that he dubs the ‘Aim Principle’ and the ‘Communication Principle’. Each of the two principles identifies a certain type of act as one by which a person can place herself under an obligation. Slightly simplifying by setting aside a proviso about validity conditions, and partly unpacking the Communication Principle so as to highlight the structure that makes it vulnerable to Lewinsohn’s critique, the two principles tell us that:

(AP) For any agent *S*, any action  $\varphi$ , and some obligations *O*: If *S*  $\varphi$ s with the intention of thereby incurring obligation *O*, then *S* is under that obligation.

(CP) For any agent *S*, action  $\varphi$ , hearer *H*, and some obligations *O*: If *S*  $\varphi$ s with the intention of getting to *H* to believe that *S* intends to thereby incur obligation *O*, then *S* is under that obligation.

Each of the two principles encapsulates a conception of the ‘fiat’ by which persons may shape the deontic situation in different respects – incurring an obligation being just one such respect.<sup>5</sup> According to the Aim Principle, the fiat consists in an act performed with a certain kind of (deontic) aim. According to the Communication Principle, the fiat consists in an act whereby a person intentionally communicates a certain kind of content (namely that she is thereby pursuing a certain deontic aim). If Lewinsohn’s argument succeeds, neither kind of act is ‘naturally intelligible’. Thus if his argument succeeds, there can be normative powers, exercisable by and among rational and informed agents, only to the extent that their exercise does not consist simply in the performance of ‘fiat’ acts as characterised by the two principles, and as being normatively significant only in the ways provided for by these principles.<sup>6</sup>

---

<sup>3</sup>In this paper I will use the following expressions interchangeably, since it does not seem to me that any part of Lewinsohn’s argument turns on distinctions that might be captured by more specific ways of using them: ‘ $\varphi$ -ing for the sake of bringing about *X*’; ‘ $\varphi$ -ing with the goal, or aim, of bringing about *X*’; ‘ $\varphi$ -ing with the intention of bringing about *X*’; ‘ $\varphi$ -ing for the purpose of bringing about *X*’; ‘ $\varphi$ -ing in order to bring about *X*’; ‘ $\varphi$ -ing so as to bring about *X*’.

<sup>4</sup>The specification relation is defined by Lewinsohn on p. 16 of his paper; more on this below.

<sup>5</sup>In what follows, I will sometimes phrase my comments in the more general register of effecting a deontic change, rather than speak specifically of incurring an obligation. This is in the spirit of Lewinsohn’s paper, which treats promissory fiats as a special case of a more general problem.

<sup>6</sup>Could a convention-independent normative power, and hence (if Lewinsohn’s argument succeeds) one whose exercise is not an act of fiat, satisfy the ordinary concept of the power to promise? Lewinsohn does not commit himself on this question. A negative answer could be prompted by the thought that it is a conceptual truth about promissory obligation that it results from fiat, and that the concept of fiat (in the relevant sense) is exhausted by acts that are vulnerable to Lewinsohn’s argument. Or it may be that normative reflection shows that obligations that have all the typical properties of promissory obligation cannot or are unlikely to result from any act that is not an act of fiat.

## 2. The argument against the Aim Principle

The Communication Principle, or something close to it, is relied on – often just implicitly – by many philosophers writing about promising and other normative powers. Reliance on the Aim Principle is much less common (as Lewinsohn mentions, its most explicit endorsement is found in debates surrounding consent). But since Lewinsohn’s argument against the former is an extension of his argument against the latter, I will follow him in focusing part of my discussion on the Aim Principle. Lewinsohn’s argument that the antecedent of the Aim Principle cannot be satisfied under conditions of full information and rationality runs as follows:

(1) The antecedent of (AP) is satisfied just in case S does something or other with the aim or intention of thereby incurring some specific obligation O. The Contribution Condition tells us that if S  $\varphi$ s with this aim, her  $\varphi$ -ing (here and now) must be something of which she believes that it could contribute to her thereby incurring O. Since we are restricting the discussion to rational and suitably informed agents, we are asking whether that belief could be true.

(2) As long as we restrict our view to those cases where S’s  $\varphi$ -ing does not contribute to her incurring O in any other way than by triggering the Aim Principle, S’s  $\varphi$ -ing (here and now) could contribute to her thereby incurring O only if it is possible that her  $\varphi$ -ing (here and now) ‘comes to account for’ her  $\varphi$ -ing with the intention of thereby incurring O. This is because of the Explanatory Constraint. In all cases in which S incurs O simply by satisfying the antecedent of the Aim Principle, her incurring O is explained or ‘accounted for’ directly, fully, and non-redundantly by her  $\varphi$ -ing with the intention of thereby incurring O. And this means that her  $\varphi$ -ing (here and now) can be part of an explanation of her incurring O only if her  $\varphi$ -ing is part of an explanation of her  $\varphi$ -ing *with the intention of thereby incurring O*.

(3) But, Lewinsohn says, it is not in fact possible that S’s  $\varphi$ -ing forms part of an explanation of her  $\varphi$ -ing with the intention of thereby incurring O. This is because of the Specification Constraint. ‘ $\varphi$ s in order to bring about X’ is a specification, in the relevant sense, of ‘ $\varphi$ s’; and that a person’s action satisfies the less specific description of a pair of descriptions related in this way cannot help explain the fact that her action satisfies the more specific one.

To sum up: If (by the Specification Constraint) S’s  $\varphi$ -ing cannot even partly account for her  $\varphi$ -ing with some particular goal, then (by the Explanatory Constraint) it cannot even partly account for her incurring O, and thus (by the Contribution Condition) S cannot  $\varphi$  with the goal of incurring O. But  $\varphi$ -ing with the goal of incurring O is what it would take to satisfy the antecedent of the Aim Principle, which therefore cannot be done.

## 3. Three concerns

It is a strikingly intuitive and satisfying argument. But there are two respects in which I find myself resisting its premises (3.1. and 3.2.). In addition, even if the argument should succeed against the Aim Principle, I am not convinced that it also succeeds against the Communication Principle (3.3.).

### 3.1. Contributing here and now

Is it true that a person can (rationally) act with the goal of bringing about some outcome X only if she believes that her so acting (here and now) could contribute to the

occurrence of X? It seems to me that at least something very close to the Contribution Condition must be true. One might in fact think that the Contribution Condition can be derived from a widely accepted belief condition on intention.  $\varphi$ -ing with the intention or purpose of accomplishing X seems to entail intending or aiming to *accomplish X by  $\varphi$ -ing* (at least among other things). For all  $\Psi$ , intending to  $\Psi$  (whether concurrently or for the future) entails believing that it is not impossible that one  $\Psi$ s (or at least not believing that it is impossible that one  $\Psi$ s). Now let  $\Psi$  stand for ‘accomplish X by  $\varphi$ -ing’. Then intending to accomplish X by  $\varphi$ -ing entails believing that it is not impossible (or at least: not believing that it is impossible) that one will accomplish X by  $\varphi$ -ing. And this commits one to believing that one’s  $\varphi$ -ing could contribute to accomplishing X (or at least to not believing that one’s  $\varphi$ -ing could not contribute to accomplishing X).<sup>7</sup>

Yet as Lewinsohn mentions<sup>8</sup>, one might wonder whether we should not view the Contribution Condition as just an approximation to a more liberal condition, one which would not support the argument against the Aim Principle. Perhaps, in order for me to (rationally)  $\varphi$  for the sake of some goal X, I need not in fact believe that my  $\varphi$ -ing, as such – regardless of the purpose with which I  $\varphi$  – could contribute to X. Instead, perhaps I need only believe that my  $\varphi$ -ing *with that aim*, i.e., *for the sake of bringing about X*, could contribute to X.<sup>9</sup>

Of course we will still want to allow, as the original Contribution Condition does, that there are goals to which one’s  $\varphi$ -ing can contribute independently of the intention with which one  $\varphi$ s. A replacement for the original Condition would therefore have to be one that can be satisfied in either of these two ways – by believing that one’s  $\varphi$ -ing could contribute to X, or by believing that one’s  $\varphi$ -ing with goal X could do so. Lewinsohn, as he must for the purposes of his argument against the Aim Principle, rejects the idea of dropping the Contribution Condition in favour of this ‘Disjunctive Condition’. I find myself doubting that this rejection is well motivated, as I will now try to explain.

I will proceed on the assumption that the Specification Constraint is true and that it entails that my  $\varphi$ -ing cannot contribute to my  $\varphi$ -ing for the sake of X. Thus replacing the Contribution Condition by the Disjunctive Condition does not depend on dropping or weakening the Specification Constraint. It does not involve the claim that my  $\varphi$ -ing could, after all, help account for something that the Specification Constraint tells us it cannot even partly account for (namely, my  $\varphi$ -ing for the sake of X). Rather, the point is that the Disjunctive Condition can be met even where my  $\varphi$ -ing, as such, in no way

---

<sup>7</sup>As stated by Lewinsohn, the Contribution Condition requires the presence of a belief in the possibility of contribution, rather than just (as the weakest version of a belief requirement on intention would suggest) the absence of a belief in the impossibility of contribution. This could appear to offer a way of escaping the grip of his argument: if no particular instrumental or ‘contribution’-type belief is required on the part of someone who  $\varphi$ s with the intention of accomplishing goal X, then a person could  $\varphi$  for the sake of <bringing about X by  $\varphi$ -ing with that intention>, even if she lacks the belief that her (simply)  $\varphi$ -ing could contribute to her  $\varphi$ -ing with that intention – as long as she does not believe that it could *not* contribute.

But I do not think that this is a possible way out, since Lewinsohn’s argument assumes conditions of full (relevant) information. In that context, the Contribution Condition can be replaced by an Objective Contribution Condition: a person who has all the relevant information can rationally  $\varphi$  with the goal of bringing about some outcome X only where her  $\varphi$ -ing under the circumstances could contribute to the occurrence of X. This Objective Contribution Condition requires not a specific belief on the part of a (rational) agent but requires rather that the content of the relevant belief be true. The Objective Condition states what needs to be true of the proposed action, rather than what needs to be true of an agent’s psychology.

<sup>8</sup>Footnote 42 of his paper.

<sup>9</sup>Stated in the objective register: If S is rational and suitably informed, S can  $\varphi$  with the aim of bringing about outcome X only if S’s  $\varphi$ -ing with that aim could contribute to outcome X.

helps contribute to realising the goal for whose sake I  $\varphi$ . To satisfy the Disjunctive Condition, it suffices that my  $\varphi$ -ing for the sake of  $X$  (which is not, we are assuming for now, decomposable into my  $\varphi$ -ing plus some further feature of the situation) be able to contribute to  $X$ .

If we replaced the Contribution Condition by the Disjunctive Condition, the argument for the impossibility of satisfying the Aim Principle and the Communication Principle would be blocked. But is the Disjunctive Condition a plausible substitute for the Contribution Condition? Lewinsohn does not think that it is. Before addressing his grounds for rejecting the Disjunctive Condition, let me highlight its general appeal.

It seems to me that a principled rationale for the Disjunctive Condition can be extracted from the clause ‘here and now’ or ‘on this occasion’ that forms part of the original Contribution Condition. It is important to see that the Contribution Condition does not require that my  $\varphi$ -ing be able (as I believe) to contribute to my goal even if circumstances at the time of acting were different from what they in fact are. The Contribution Condition requires only that my  $\varphi$ -ing be able (as I believe) to contribute to my goal *under the present circumstances*. And that seems right. Few things I do for the sake of some goal are capable of contributing to that goal no matter the further circumstances. Usually the world (including my own future self) needs to play along in some way or other. We may then ask which aspects of an agent’s situation are included in these present and future circumstances relative to which the agent must believe her action is capable of contributing to her goal. The most natural answer appears to be: all those aspects of the agent’s situation that she believes to obtain<sup>10</sup> (especially including her own future intentions, and perhaps also her own present intentions).

Consider the following example: Could pouring myself a glass of water come to contribute to quenching my thirst? Not unless I proceed to raise the glass to my mouth and drink from it. But that, in turn, will (let us suppose) only happen if I have – at the time of drinking, if not yet at the time of pouring – the aim of taking a sip of water. Absent that (later) aim, at least according to plausible ways of filling in the story, my pouring will not contribute to my drinking. Thus I could not pour a glass with the aim of quenching my thirst if the Contribution Condition required that the initial action be able to contribute to that goal regardless of my own future intentions. This illustrates the fact that an entirely context-invariant interpretation of the Contribution Condition would impose a requirement so strict that it would make many ordinary actions impossible. But if, as in the example just given, my own future intentions can be among the aspects of my situation in virtue of which the Contribution Condition is satisfied in a given case, then why not my own present intentions too?

The answer to the last question depends on whether it is plausible to think of the fact *that an agent’s  $\varphi$ -ing is, at least in part, directed at (or guided by the aim or intention of) realising a specific purpose  $X$*  as part of the circumstances of her  $\varphi$ -ing. Is this a fact whose presence I could rely on, without irrationality, in the same way in which I rely on other present or future circumstances that need to be in place in order for my action to be able to contribute to the goal at which it is aimed? If it is, then the Disjunctive Condition is just a way of making explicit what is already contained in the Contribution Condition as

---

<sup>10</sup>Or, under idealised conditions of full information, as suitable for the purposes of Lewinsohn’s argument: all those aspects of which the agent can have knowledge.

originally stated. (If it is not, the Disjunctive Condition could still be true in its own right.)

Lewinsohn rejects the idea that the Disjunctive Condition could be an implication of or a substitute for the Contribution Condition. He argues that the particular way of motivating the Disjunctive Condition that I have just sketched would provide equal support for an *even more liberal* condition, of which the Disjunctive Condition would only be a special case, and that this More General Condition (as he terms it) ‘fails on extensional grounds’. With the ‘here and now’ understood in the inclusive way I have suggested, the Contribution Condition would be satisfied by *any* specification of  $\varphi$  that S will make true (as she believes) if S  $\varphi$ s on the occasion. But settling for this More General Condition, Lewinsohn thinks, has absurd implications.<sup>11</sup> And that is meant to impugn any construal or replacement of the Contribution Condition that turns on the suggested inclusive reading of the clause ‘here and now’ or ‘under the present circumstances’.

I am not persuaded that the More General Condition has the false implication that Lewinsohn argues it does, but I will only register my doubt in a footnote here.<sup>12</sup> Supposing instead for the sake of the argument that the More General Condition does in fact imply a falsehood, this would be a reason to reject it but not, it seems to me, a reason to reject the Disjunctive Condition (assuming that the latter does not have the false implication). If so, it turns out not to be the case that, as Lewinsohn seems to suggest, we don’t have sufficient reason to prefer the Disjunctive Condition over the More General Condition. The heart of his claim here is that there is no *principled* reason to prefer the Disjunctive Condition. But that the Disjunctive Condition is not vulnerable to a certain counter-example, while still explaining what needs to be explained, seems to me to be as good a reason as any to prefer it to a condition that is. Perhaps the problem is that with the lapse of the More General Condition, we would lack a principled reason to prefer the Disjunctive Condition over the Contribution Condition understood in the original narrower way. That would return us to the imagined charge that the Disjunctive Condition is somehow objectionably *ad hoc*.

I am not sure, in general, how to adjudicate the question whether a proposed principle is *ad hoc*, or what to say about the dialectical force of such a diagnosis where it is true. So I

---

<sup>11</sup>Lewinsohn, footnote 42.

<sup>12</sup>Lewinsohn points out, in effect, that the More General Condition is consistent with the following conditional: If S performs some action A\* (e.g., taking a red pill) in order to achieve goal G, in the belief that doing A\* can contribute to G, then for any action A (e.g., taking a pill of any colour) of which A\* is a specification that will be realised if S does A under the circumstances, S can also be said to be doing A in order to achieve G (even where S does not believe that all ways of doing A could contribute to G). And that conditional seems false.

But that the More General Condition is consistent with this conditional does not show that the More General Condition commits us to accepting it. After all, the More General Condition, like the Contribution Condition, articulates only a necessary condition for  $\varphi$ -ing with some goal G, not a sufficient condition. It seems plausible that at least one further necessary condition is (i) that S believes that she is  $\varphi$ -ing (e.g., doing A / taking a pill of any colour), and perhaps also (ii) that she believes that she is  $\varphi$ -ing with goal G (e.g., that she has the belief: ‘I am taking a pill of any colour in order to learn the truth’). The absurdity, such as it is, of saying that I took a pill of any colour (as opposed to: a pill of *some* colour) in order to learn the truth may be due to the fact that on the most natural reading of that sentence, one or both of these further conditions are not met. That the More General Condition does not by itself do the work done by these further conditions does not show that it is extensionally inadequate.

Such possible further conditions aside, that I took a pill of any colour in order to learn the truth is naturally read as implying that in taking the pill, I was indifferent to its colour; and this is in tension with the More General Condition itself, which after all requires that I believed that I was taking the red pill (a belief that would in standard cases have been founded on my intending specifically to take the red pill, as opposed to being indifferent between the act of taking the red pill and the act of taking the blue pill).



will not further pursue those questions here. In any case they would need answering, in this context, only for those who are satisfied that Lewinsohn's counter-example is in fact an implication of the More General Condition.

### 3.2. Weakening the specification constraint

The Specification Constraint claims that where two act descriptions stand in the 'specification' relation defined by Lewinsohn, the fact that an action satisfying the more general description is performed could never contribute to (or at least partly 'account for') the fact that an action satisfying the more specific description is performed. But as in the case of the Contribution Condition, I wonder whether here too there may be a weaker and no less plausible alternative constraint in the vicinity.

Perhaps the fact that an action satisfies a less specific description cannot (even partly) account for the fact that the action satisfies a more specific description *unless* by satisfying the less specific description *here and now, i.e., under the present circumstances*, one's action is guaranteed to satisfy the more specific description. In the sort of case under consideration, might we not think that when an agent has the general goal of incurring obligation *O* in some way or other (a possibility that Lewinsohn does not dispute), and the agent is able to perform an action  $\varphi$  of which she knows that performing the action *with that goal* could contribute to her incurring *O*, then if she  $\varphi$ s at all it is likely or even certain that she will be  $\varphi$ -ing with that goal? And if that is so and she knows that it is, could she then not  $\varphi$  in order to bring about a goal that (as she knows) can be brought about only by  $\varphi$ -ing with that very goal?

Compare cases like the following: Couldn't someone choose to (simply) read a book in order to be reading a book distractedly, or (simply) lie in a hammock in order to lie in it lazily, or (simply) listen to some music in order to listen to it incomprehensibly – provided she knows that under the circumstances, if she reads or lies or listens at all, she will respectively be doing so, through no further choice of hers, distractedly, lazily, or incomprehensibly? Could someone not go down a slide in order to go down a slide fast, provided she knows that if she does go down this slide now, she will be going fast? Or could someone not disclose a particular fact so as to disclose that fact selfishly, if she knows that under the circumstances, disclosing it will be a selfish act? I am supposing that in these and similar cases, at least under some interpretations of them that are consistent with their possibility, the second description is a specification of the first in the sense defined by Lewinsohn.

As Lewinsohn points out, the source of the Specification Constraint is the Contribution Condition. And the weakening of the Specification Constraint that I am here proposing picks up, once again, from the 'here and now' or 'on the occasion' clause contained in the Contribution Condition. (The clause tends to drop out of view in Lewinsohn's discussion of the Specification Constraint, but it belongs to it by right.) The fact that there is an action satisfying a more general description, Lewinsohn says, cannot go any way towards explaining the fact that there is an action satisfying a more specific description when the two actions stand in the specification relation to each other. But that does not seem so clear in cases where performing an action that satisfies the more general description will ensure, under the circumstances, that the action also satisfies

the more specific description.<sup>13</sup> In those cases, it seems that performing an action satisfying the more general description can in fact contribute to, or partly account for, having performed an action that satisfies the more specific description.<sup>14</sup>

In the case under discussion, i.e., with respect to the Aim Principle, what we are asking about is the possibility of doing something in order to be doing that very thing with a specific aim or purpose or intention (and more remotely, in order to incur an obligation that can only be incurred by doing that thing with that intention). Could I  $\varphi$  in order to  $\varphi$  for the sake of goal G? My suggestion is that I could: My  $\varphi$ -ing could partly account for, and thus contribute to, my  $\varphi$ -ing for the sake of G, *provided* that (as I know) under the circumstances, if I  $\varphi$  at all my  $\varphi$ -ing will be guided by G.<sup>15</sup> It may be that there is some special difficulty when the goal in question is one which, as I know, can be realised only by acting for the sake of that very goal. But whatever the nature of this difficulty, it does not seem to derive from the Specification Constraint.

### 3.3. A modified communication principle

Lewinsohn argues that like the antecedent of the Aim Principle, the antecedent of the Communication Principle cannot be satisfied. It is not possible for a person S to perform any action  $\varphi$  (standardly, an utterance) with the ‘communicative aim’ of getting someone to believe that S intends (i.e., has the ‘communicated aim’) to *thereby* incur an obligation O. It is not possible to do this because it is not possible, under conditions of rationality and full information, to aim to get a hearer to believe something that could not be true; and it could not be true that S has the ‘communicated aim’: the aim of incurring an obligation O in the way identified by ‘thereby’.

Why could it not be true that S has the communicated aim? As stated, Lewinsohn’s argument against the Communication Principle proceeds on the assumption that

<sup>13</sup>Can we transform Lewinsohn’s own example – killing, versus killing by stabbing – into a case of this type? Perhaps so: Suppose that the killer, at least on this occasion – or perhaps on any occasion –, has no way of killing other than by stabbing, and she knows this. In that case, she knows that if she kills the victim at all, she kills him by stabbing him. I grant that it remains counter-intuitive to say that she can therefore kill him with the aim of killing him by stabbing him. But the source of that strangeness, I think, is not the specification constraint. I will not try to give a full account of what the source may be, but I suspect that it may have to do with the fact that we imagine the stabbing to be done, in this case, with the aim of killing, and that a person cannot do A (killing) in order to <do B in order to do A>. But in fact killing (simpliciter) need not be an aim the person has in this case. Her only aim may be specifically to kill by stabbing. In order to realise that aim, it is not sufficient that she stab; she must also (thereby) kill. It is true that under normal conditions, killing, as such, does not bring one closer to having killed by stabbing. But in a situation in which killing x necessarily consists in stabbing x, it does (and not just closer but indeed all the way).

<sup>14</sup>If I were to try and pinpoint where in Lewinsohn’s chain of reasoning this possibility gets ruled out, I would say that it is in his taking a certain restrictive definition of the relation of ‘partially explaining’ to be the only alternative to the relation of fully explaining or the absence of any explanatory relation (Lewinsohn, p. 16 f.). It seems to me that some fact A can figure in an explanation of another fact B even where it is neither the case that A fully explains B, nor that A is a non-redundant part of some set of facts that fully explains B (as would be required for ‘partial explanation’ under Lewinsohn’s definition). I suspect that if my examples are sound, they show that a fact can sometimes be explanatory even where it is a redundant part of a fully explanatory set. That A, B, and C are jointly sufficient for explaining X does not entail that a further fact D (of which one among A, B, and C is a specification) could not also contribute to an explanation of X, as long as the presence of D – under the circumstances – entails the presence of A, B, or C.

<sup>15</sup>What might make this belief true? One possible reason, not necessarily the only one, is that G (incurring an obligation) is a goal that I have prior to and independently of my considering exactly how to realise G, and I believe that  $\varphi$ -ing, on the occasion, would contribute to realising G, and I have the options of (i) not- $\varphi$ -ing or of (ii) setting aside G in considering whether and how to  $\varphi$ , and I refrain from taking those options. To believe that  $\varphi$ -ing, on the occasion, would contribute to realising G, it seems sufficient that I believe the Aim Principle and that I accept only the relaxed version of the Specification Constraint.

‘thereby’ here means ‘by making some utterance U’.<sup>16</sup> Given that assumption, the argument for the impossibility of satisfying the Aim Principle can be run against the possibility of making some utterance U in order to incur an obligation (simply in virtue of triggering the Communication Principle). S cannot have this communicated aim, because – on this reading of ‘thereby’ – having that aim would consist in aiming to trigger the Communication Principle (= to perform an action with the communicative aim) simply by  $\varphi$ -ing (i.e., simply by uttering U); and this, in turn, would require that S’s  $\varphi$ -ing (uttering U) could contribute to her  *$\varphi$ -ing with a specific aim C*. Since  *$\varphi$ -ing with aim C* is a specification of  *$\varphi$ -ing*, such a contribution is ruled out by the Specification Constraint. Knowing this, S cannot have the communicated aim, and thus could not intend to get a (rational and informed) person to believe that she has it. Therefore she cannot, among rational and informed persons, make any utterance with that intention, and thus she cannot satisfy the antecedent of the Communication Principle.

But on an alternative reading of the Communication Principle, it does not seem to be facing this problem. Let us take ‘thereby’ (in the communicated intention) to refer not simply to S’s  $\varphi$ -ing (= uttering U), but instead to her  *$\varphi$ -ing with the communicative aim* (that is, the aim of getting H to believe that she intends to incur O by  $\varphi$ -ing with that aim). What S then needs to communicate is her intention of incurring O not by  *$\varphi$ -ing (uttering U) simpliciter*, but – irreducibly – by  *$\varphi$ -ing / uttering U with the communicative aim*. What it takes for the Communication Principle to be triggered, on this reading, is merely that a person should make an utterance with the goal of conveying an intention to *<incur O by making an utterance with the aim of conveying that intention>* (where the italicised phrase stands for the revised interpretation of ‘thereby’). This does not fall foul of the Specification Constraint.

There is a distinct resemblance between this proposal and my remarks on the Contribution Condition and the Specification Constraint. But the current proposal is consistent with accepting both of those conditions in their original form.<sup>17</sup> The proposal concerns instead the content of the normative principle under discussion. The antecedent of the modified Communication Principle can be satisfied even where it is known to all parties that uttering some sentence U could not contribute to uttering U with the communicative intention, and thus that uttering U could not be done with the aim of uttering U with the communicative intention. All it takes to trigger the modified principle is to utter U with the communicative intention. So all that the hearer needs to be able to

<sup>16</sup>This assumption is implied by the letter of his argument on pp. 22 f. But as Lewinsohn has pointed out to me, his claim (in footnote 41 of his paper) that it does not ultimately matter to his argument how we resolve the ambiguity with respect to ‘thereby’ is meant to apply not just to the argument against the Aim Principle but equally to the argument against the Communication Principle. While I agree that the resolution of the ambiguity does not bear on the argument against the Aim Principle, it seems to me that the same does not hold with respect to the Communication Principle. More on this below.

<sup>17</sup>The Contribution Condition is not in question: To have the aim of bringing about some result X (in this case: incurring obligation O) by  $\varphi$ -ing with intention C (in this case: by uttering U with the intention of getting the hearer to believe that I intend to bring about a normative change in a certain way), what I need to believe, by the Contribution Condition, is that my  $\varphi$ -ing with intention C could contribute to X’s coming to pass. There appears to be no problem in believing this, if the modified Communication Principle is true and believed by me.

The Specification Constraint is not in question either. It states that I cannot believe that my  $\varphi$ -ing (uttering U) could even partly account for my uttering U with the communicative intention; and the modified Communication Principle can be satisfied consistently with this. Satisfying its antecedent does not require that my aim in uttering U is to do the thing that would fully and directly and uniquely account for my incurring obligation O, since satisfying its antecedent does not require that I believe that I could incur obligation O (partly) by uttering U – only that I could incur O by uttering U with a suitable intention.

believe about the speaker is that the speaker has the aim of incurring, and thus the belief that she could incur, an obligation O by <uttering U with the communicative intention>.

One might challenge the proposal just made by asking: Is it in fact possible to do something,  $\varphi$ , with some aim A so as to thereby – i.e., by doing it specifically with that aim – realise some further aim B? In the case at hand: Could one make an utterance with the aim of getting someone to have some belief B, so as to thereby incur an obligation O? There may be difficulties in understanding this sort of structure, perhaps ones that will end up sinking the modified Communication Principle. I am inclined to think that at least some of what tends to be perceived as puzzling about such cases can be made good sense of. Making sense of it would go some way towards vindicating the possibility of what Joseph Raz has called positive second-order reasons.<sup>18</sup> But whatever difficulties there may be here, it does not seem that they are identical with the ones asserted by Lewinsohn's argument against the Aim Principle.

#### 4. Normative powers and social rules

Lewinsohn aims to show that we cannot effect normative changes by triggering the Aim Principle or the Communication Principle, or any like principles for normative changes other than incurring an obligation. But of course he does not deny that we are in fact able to effect normative changes, and that we can do so both intentionally and with the intention of effecting such changes. Nor is he denying, I think, that among the things we can intelligibly do with the intention of effecting particular normative changes are ones that do not presuppose any particular social rules, at least not social rules specifically conferring the power to effect these changes.

So what exactly is it that creates the pressure to think that power-conferring social rules must have an important role to play in any satisfactory account of our ability to change the normative situation? Lewinsohn's thesis is not just that *there are* some such social rules (of which some legal rules may be an instance), and that there are some ways in which people can intentionally change the normative situation by triggering such rules. It may be that not many would disagree with that view, but it is not clear how it would derive any specific support from the unintelligibility argument advanced in the first parts of Lewinsohn's paper, or require that support. So what exactly is it that power-conferring social rules are needed to explain?

The positive thesis of the paper – call it the conventionalist thesis – seems to be that there are *some particular ways* of changing the normative situation such that we can effect normative changes in *those ways* only in virtue of the existence of power-conferring social rules. But which ways are these? We cannot identify them as simply those that involve power-conferring social rules, on pains of rendering the thesis trivial. Lewinsohn's definition of a power-conferring social rule, loosely modelled on a proposal about power-conferring legal rules by Raz, points to a non-trivial answer: the relevant ways are those of which it is true that people by and large produce normative changes in those specific ways only if they intend to do so. It is *those ways* of producing normative changes – in Lewinsohn's terminology, intent-screening ways – that can be explained only by the existence of suitable social rules.

---

<sup>18</sup>Cf. Joseph Raz, *Practical Reason and Norms* (Oxford University Press 1999), Chapter 1.2 and postscript to the second edition.

If this is what Lewinsohn has in mind, there seems to remain conceptual space between (i) the concept of an intent-screening way of changing the normative situation that is due to the existence of a power-conferring *social* rule, and (ii) the quite specific, and putatively unintelligible, intent-screening ways characterised by the principles that Lewinsohn discusses. Supposing that those principles do in fact need to be discarded, why should we think that the only remaining intent-screening ways of changing the normative situation must be ones involving social rules? Is there some general reason to believe that a way of changing the normative situation cannot be intent-screening unless a suitable group of people believe, or at least accept for practical purposes, that this is indeed a way of changing the normative situation?<sup>19</sup> If so, I am not sure that Lewinsohn tells us what that reason is. Of course, even if it is true that he does not tell us, it does not follow that there is no such reason.

There is a different kind of thing that might be meant by ‘a social rule’: not a rule that is socially practised or accepted (among some relevant group), but one that makes reference to social facts or conventions in its content, by claiming that they have a certain normative relevance. Being a social rule in the first of these two different senses can go together with being a social rule in the second sense, but it need not. That one must never lie can be a social rule in the first sense but not in the second (and, unless we rely on some more than minimal conception of a social rule in the first sense, its content could be either true or false quite independently of any facts about whether it is believed or practised). By contrast, it would be a social rule in the second sense – even where nobody believes or accepts it – that it is always fine to lie under conditions under which most others would do so, or would think it fine to do so. And again, the content of that rule could be either true or false.<sup>20</sup>

It is clear from Lewinsohn’s discussion that he classifies rules as social rules not in virtue of their content but in virtue of the fact that they are accepted or practised. But any rule, identified just in terms of its content, could either be socially accepted or practised, or not. If a certain socially accepted rule would indeed provide for an intelligible intent-screening way of changing the normative situation, then why should that rule, identified in terms of its content, not be able to perform the same function even when it is not socially accepted? Or to put the same question differently: why should we think that a rule can be power-conferring only if it is socially practised?<sup>21</sup> Lewinsohn

---

<sup>19</sup>I am assuming that the property of being believed, accepted, and/or followed – in short, being socially practised – belongs to the minimal concept of a social rule, in one familiar sense of ‘social rule’. Some conceptions require significantly more than this: for instance in stipulating that social rules are only those socially practised rules that are binding (valid, true) partly in virtue of being socially practised; or only those whose general acceptance is itself part of people’s reason for accepting the rule. Lewinsohn explicitly sets aside the question whether the social rules at issue in his discussion are morally binding or what the conditions of their being morally binding are. Nor does he take any stance on whether for a rule to be a social rule its acceptance by others must be among the reasons for people’s accepting it.

<sup>20</sup>If its truth depends on whether it is socially practised, the rule would be a social rule in yet a further sense that we may distinguish from the first two: a rule that is true or valid only if practised, or for that matter (and there may be no actual instances of this): true or valid only if not practised. A rule could be a social rule in this third sense, in having social acceptance or non-acceptance among its validity conditions, even when it is neither practised in any community nor contains any reference to social facts. The rule ‘One must never lie’ would be a social rule in this third sense if, for instance, the validity of the rule (perhaps restricted to some context) depended on whether it was widely accepted (in that context). Those who reject social relativism will explain any rule belonging to the third type by reference to some rule of the second type.

<sup>21</sup>There may be much to be said for accepting a weaker requirement: a rule that can make our action intelligible must be one that is conceptually and epistemically accessible to all relevant parties, and perhaps one that all parties have sufficient reason to believe. But it would seem to be stretching the concept of a social rule to say that all rules meeting these conditions are therefore social rules.

puts very significant pressure, at the least, on a prime candidate for a type of power-conferring rule that could be valid independently of being socially practised (namely, the Communication Principle). But it is not clear to me why ruling out this candidate should go any way towards an impossibility argument. Conversely, that the absence of social acceptance is not the key problem is suggested by the fact that, as Lewinsohn points out, the unintelligibility argument applies to the principles indicted by him whether or not they enjoy such acceptance.

## 5. Conclusion

All of the points examined here would deserve further discussion. There are also other facets of Lewinsohn's exceptionally rich paper that would deserve engagement in their own right, among them the illuminating discussion of intent requirements – including possible attempts to immunise the targeted principles against the unintelligibility argument by appealing to objective intent – and Lewinsohn's remarks on the Humean origins of his argument. My primary concern has been to bring out some ways in which one might attempt to extricate what Lewinsohn rightly characterises as the prevalent non-conventionalist account of normative power from the grip of the unintelligibility charge. Whether or not it is possible to defend that account, his paper prompts us to give serious thought to alternatives, including though perhaps not limited to ones that look to socially accepted rules.<sup>22</sup>

## Disclosure statement

No potential conflict of interest was reported by the author.

## Funding

Work on this paper was supported by a grant from the Swiss National Science Foundation (grant no. 200749).

## ORCID

Felix Koch  <http://orcid.org/0009-0002-5047-2033>

---

<sup>22</sup>In memory of Joseph Raz, and with gratitude. I would like to thank Ulrike Heuer, John Hyman, and Yuuki Ohta for convening the workshop 'Reasons and Normativity: Themes from the Philosophy of Joseph Raz'; Benjamin Kiesewetter for his reaction to a draft of this article; and Jed Lewinsohn for his captivating and wide-ranging paper and for rewarding exchanges about a number of the points discussed here.