Resolutions provide reasons or: "how the Cookie Monster quit cookies"

Adam Bales[1] · Toby Handfield[2]

© The Author(s) 2021

Abstract

Why should we typically act in accordance with our resolutions when faced with the temptation to do otherwise? A much-maligned view suggests that we should do so because resolutions themselves provide us with reasons for action. We defend a version of this view, on which resolutions provide *second-order* reasons. This account avoids the objections typically taken to be fatal for the view that resolutions are reasons, including the prominent bootstrapping objections.

Keywords Resolutions · Temptation · Reasons · Bootstrapping arguments

The Cookie Monster wants to cut back on sugar. However, when faced with cookies he feels tempted to consume and experiences motivated cognition: his desire leads his beliefs to shift, so that he temporarily believes he will be able to stop eating after just one. This belief is false. Consequently, he eats not just one, but several cookies. Hoping to overcome temptation, the Cookie Monster resolves to eat cookies only on weekends.

On Monday, the Cookie Monster encounters cookies and experiences temptation. As a result, he strongly desires cookies and believes he can stop after eating just one. These beliefs and desires seem to entail that he should ignore his resolution and consume (as he now believes that he can satisfy his desires without health costs). Yet this is absurd: surely the Cookie Monster may rationally stand by his resolution, given that it was formed to help him resist this very temptation. So we have a puzzle: what

✉ Adam Bales
   atb39@cam.ac.uk

1   University of Cambridge, Cambridge, UK

2   Monash University, Melbourne, Australia

appears rational, given his beliefs and desires, appears irrational, given his resolution. Further, as desire and judgement shift often accompanies temptation, this is no parochial problem: it fundamentally threatens the normative force of resolutions.[1]

Fortunately, there's a simple solution here. Imagine that the Cookie Monster resists and we ask why. The Cookie Monster might reply that he resisted because he resolved to do so. Taken at face value, this suggests that a resolution to ϕ provides a reason to ϕ. So the Cookie Monster's reasons to refrain were not just the health benefits *but also the fact that he resolved to refrain*. These reasons can then justify standing by the resolution. Unfortunately, there's something of a consensus that this account fails (cf. Bratman 1987; Broome 2001; Holton 2009). However, in this paper we will defend the account. In particular, in the first half of the paper, we will argue that resolutions provide reasons for action, without discussing the nature of these reasons in detail. Then, in the second part of the paper, we will clarify the view, by suggesting that the sorts of reasons provided by resolutions are *second-order* reasons, where such reasons operate by changing the normative relevance of standard first-order reasons.

1 Reasons generation

Our interest is in a subjective notion of rationality, on which what one ought to do depends on one's mental states. So the above puzzle flows from the fact that the Cookie Monster's beliefs and desires seem to entail that he *subjectively* ought to eat the cookies. Likewise, we will argue that resolutions provide *subjective* reasons.[2]

Keeping this in mind, the basic reasons provision account is simple: resolving to ϕ provides an agent with a reason to ϕ. However, two issues arise. First, foolishly-formed resolutions lack normative force. So, at most, only *rationally-formed* resolutions provide reasons. Second, it is sometimes rational to abandon a rationally-formed resolution. For example, if I resolve to jog, I should abandon this resolution if I break

---

[1] The challenge faced by the Cookie Monster is sequential: it results from the fact that he faces a sequence of opportunities to consume cookies. In contrast, some other cases where resolutions are helpful involve one-off choices. Still, this distinction is unimportant for our purposes: what is needed is a general account of why it's rational to act in accordance with our resolutions (whether or not the case is sequential).

Note that there are various ways we might interpret the Cookie Monster case. Perhaps the Cookie Monster suffers from motivated cognition. Perhaps he experiences preference change. Perhaps his case is similar to that facing Quinn's self-torturer (Quinn 1990). Depending on the interpretation, the Cookie Monster might be able to avoid problematic behaviour without forming a resolution. For example, if the problem results from motivated cognition then perhaps cognitive therapy would allow the Cookie Monster to better control his thoughts and so resist cookies without a resolution. We acknowledge this possibility but don't find it concerning. For a start, there's still a question of what behaviour is rational for an agent who can't access these alternative solutions (what if cognitive therapy isn't available?). Further, we think that resolutions can remain a valuable and interesting normative tool, even if the benefits they offer can also be achieved in other ways.

[2] Those who reserve talk of "reasons" for objective contexts should read our talk of reasons as talk of considerations that partially determine an agent's subjective obligations. Note that some of the literature to be discussed engages with more objective questions than ours. Still, such literature can often provide insight into the subjective question.

both legs. So a resolution has normative force only when it is rational to *retain* that resolution.[3]

We are led to:

RESOLUTIONS- PROVIDE- REASONS (RPR): A rationally-formed, rationally-retained resolution to ϕ provides an agent with a reason to ϕ.

This is the account that we will discuss.[4]

2 The pragmatic argument

Pragmatic arguments can be provided for RPR. Before outlining such an argument, we first need to distinguish normative reasons from motivating reasons. Normative reasons, which we simply refer to as reasons in this paper, provide normative support for some action. So, at least on a simple account, if the balance of normative reasons support some action then we ought to carry out that action. On the other hand, motivating reasons motivate an agent to carry out some action. Such reasons are subjective: the presence of chocolate in the fridge may motivate some people to approach the fridge but others to avoid it. Now, it's clear that the mere fact that something provides a motivating reason doesn't mean it provides a normative reason: some agents may be irrational in their motivations and so such motivations are hardly guaranteed to reflect normative considerations. Still, there might seem to be a connection between motivating reasons and normative reasons. In particular, it might seem that if an agent would end up well off if some consideration provided them with a motivating reason, then this consideration provides a normative reason. If so, then we can provide a pragmatic argument for RPR along the following lines:

BENEFIT: Pragmatic benefit would accrue to agents if (rationally formed and retained) resolutions provided sufficiently strong *motivating* reasons for action.[5]

ENTAILMENT: If BENEFIT is true then (rationally formed and retained) resolutions provide *normative* reasons for action.

BENEFIT is plausible. Consider the Cookie Monster: if resolutions provide sufficiently strong motivating reasons then his resolution helps him resist the cookies.

---

[3] In this paper, we won't provide a detailed account of when resolutions may rationally be formed and retained (though see Broome 2001, p. 113 and Holton 2009, p. 75 for discussion). Indeed, we're doubtful that any simple account of this sort can be provided, and so doubtful it could be presented merely in passing. However, it's important to note that this raises a potential challenge: it might be worried that this will make our account immune to criticism. If any counterexample is presented, we can simply arbitrarily claim that it involves a resolution that is either not rationally formed or not rationally retained and therefore isn't a true counterexample. Fortunately, even without a theory, we have a good grasp of some cases where resolutions are, or are not, rationally formed or retained. In order to avoid the arbitrariness trap, we will rely only on claims about rationality of formation and retention that we take to be extremely plausible. See footnotes 15 and 17 for specific details.

[4] See Verbeek (2014) for a related view.

[5] This formulation is deliberately vague in specifying the reliability of the connection between resolutions providing motivating reasons and pragmatic benefit. More precise formulations might say "Expected pragmatic benefit..." or "would probably accrue", or "would typically accrue", etc. For related discussion in the context of Holton's account, see Cohen and Handfield (2010), pp. 922–923.

However, if resolutions do not provide such reasons then he will plausibly eat the cookies, given his strong desire to consume and his motivated beliefs. So the Cookie Monster benefits if resolutions provide sufficiently strong motivating reasons. (Ignore any objective benefit here: we mean simply that Cookie Monster does better, subjectively, by the light of his own preferences, values, etc.)

More generally, if resolutions provide motivating reasons then they can help agents resist anticipated temptation. However, if resolutions do not provide such reasons then an agent cannot use resolutions in this way. So agents benefit if resolutions provide motivating reasons.[6] BENEFIT looks good. How about ENTAILMENT? We noted the initial plausibility of this claim above, but does it stand up to scrutiny?

2.1 Gauthier on practical reasoning

Historically, arguments for ENTAILMENT have appealed to Gauthier's (1994) claim that practical rationality is self-regulating, in that the norms of practical rationality are the norms that it is rational to choose to follow. Abstracting away from the details of Gauthier's view, we might be tempted to spell this out as:

> GAUTHIER: Something provides normative reasons for action if, and only if, it would be rational to choose for it to provide motivating reasons for action.

We can now argue for ENTAILMENT. After all, other things being equal, it is plausibly rational to make choices that pragmatically benefit us. Now, assuming all else is equal, BENEFIT entails that it's rational to choose for resolutions to provide motivating reasons. Given GAUTHIER, it follows that resolutions provide normative reasons. We have ENTAILMENT.

Unfortunately, GAUTHIER is doubtful, for reasons discussed by Velleman (1997). We won't rehearse Velleman's full discussion but a key point is that Gauthier's view involves a circularity: on this view, what we are rational to choose depends on the norms of rationality but these norms themselves depend on what we are rational to choose. This means that we cannot evaluate norms of rationality without first arbitrarily choosing some such norms. So a pragmatic argument for RPR based on GAUTHIER is problematic.[7]

2.2 An epistemic principle

Fortunately, an alternative pragmatic argument is available. First, note that there is a close, perhaps analytic, connection between practical rationality and pragmatic benefit. This is most clear if we focus on instrumental rationality; after all, to be instrumentally rational just is to carry out those actions that best promote (expected) success. However, the connection remains when we focus on a broader notion of practical rationality. Even if such a broader notion isn't solely about success, success plays a central role

---

[6] BENEFIT could be justified empirically, by reference to the actual success of agents. If the truths of rationality are knowable *a priori*, pragmatic arguments will then be heuristics (that may lead us astray if we are misled by contingent truths).

[7] Or so we will assume: this assumption makes our life harder and so hardly weakens our eventual argument.

in any plausible notion of practical rationality. Consequently, if an agent succeeds, and if this success results from her actions (rather than luck), this provides prima facie evidence that the agent acts rationally.[8] Perhaps this evidential import can be overcome by pointing to undermining considerations, but the prima facie evidence remains unless such considerations are identified.

Now, turn back to resolutions and focus on agents for whom resolutions provide motivating reasons. The above reasoning suggests that if such agents are successful (and this success results from their actions) then we have prima facie grounds to think they're rational. Further, a rational agent's motivating reasons will be the normative reasons in play.[9] So if the agents under consideration are rational then resolutions provide normative reasons. Conclusion: if agents for whom resolutions provide motivating reasons are successful (as a result of their actions) then we have prima facie evidence that resolutions provide normative reasons. This doesn't suffice to establish ENTAILMENT, which doesn't have the required *prima facie* flavour, but it does support:

> ENTAILMENT*: The truth of BENEFIT provides *prima facie* evidence that (rationally formed and maintained) resolutions provide *normative* reasons for action.[10]

That is: if it would be beneficial to be motivated by resolutions then we have apparent grounds to believe that resolutions provide normative reasons.

So now we have both BENEFIT and ENTAILMENT*. From these two principles, we reach the conclusion that we have prima facie evidence that (rationally formed and retained) resolutions provide normative reasons for action.[11] That is, we have prima facie evidence for RPR.

This argument differs crucially from the one we considered earlier. The earlier argument aspired to provide deductive proof of RPR. However, the current argument

---

[8] This informal claim can be made formally precise by appeal to probabilistic reasoning. If P(Successful | Rational) > P(Successful | Not-Rational), then upon learning that someone is successful, via Bayes rule, we should now assign a higher probability to them being rational—or at least not regard this as less probable than it was before. The degree of credence change required will be greater the rarer success is in the relevant domain, and the greater the difference in success rates between the rational and the non-rational—in other words, the less influence that luck has in that domain.

[9] An agent's motivating reasons are those things they treat, in their behaviour, as providing pro or con support for actions. Normative reasons are the things that *should* provide such support. So for a rational agent, the normative reasons just are the motivating reasons. (There may be exceptions to this link in Buridan's Ass cases, where a rational agent must choose despite a lack of reasons favouring either option. However, our appeal to cases where resolute agents are successful are not of this sort, so we ignore this exception here.) A different argument here would proceed by noting that one role of resolutions is to help agents resist temptation. As temptation is often accompanied by strong desires and motivated cognition, this means that an agent will often have (subjective) normative reasons to act in accordance with temptation. So if a resolute agent is to be rational, it must be because some other normative reasons counters the reasons favouring temptation. Plausibly, resolutions provide these reasons. (This is a form of inference to the best explanation, so taking this line would require further defence of the claim that RPR is the best explanation available).

[10] At least assuming the success of the agents results from their action. This is a safe assumption in the case at hand.

[11] There are other ways the Cookie Monster might incur this benefit, even if resolutions don't motivate him. For example, perhaps the Cookie Monster is motivated by the fact that he previously had reasons to bind himself to later action. Alternatively, the Cookie Monster might habitually act in accordance with his resolutions without further reflection (Holton 2009). So why think that considerations of success support our own view, rather than some other?

merely shows that we have prima facie evidence for the truth of this view. Still, this is enough to shift the initial burden of proof: we have evidence for RPR, so it's up the opponent of this view to show why we should reject it even so.

3 The platitude argument

This pragmatic argument can be bolstered via a platitude argument. After all, here's something we know about resolutions: it is typically rational to act in accordance with well-formed resolutions in the face of anticipated temptation. We know this in the abstract. We also know it in concrete circumstances when, for example, we evaluate someone positively if she adheres to a (sensible) diet. So our theorising should account for this platitude.

Now we can argue as follows: we should accept RPR because it makes sense of this platitude.[12] After all, given RPR it will typically be rational to respect our resolutions when facing temptation (because the resolution provides a reason to do so).

Of course, this argument is not decisive, as there are alternative ways to account for the relevant platitude (cf. Bratman 1987; Broome 2001; Holton 2009). Still, we do not aim to evaluate all possible accounts of resolutions here. Rather, we intend to show that the much-maligned reasons provision account is a serious contender.

4 Bootstrapping challenges

Unfortunately, it is widely accepted that RPR is undermined by bootstrapping arguments (cf. Bratman 1987, pp. 24–27 Broome 2001).[13] At least three such arguments exist.

4.1 Irrational bootstrapping

A simple bootstrapping argument asks us to consider some action that is clearly irrational, like stabbing ourselves or voluntarily eating parsnip (cf. Bratman 1987, pp. 24–27). Now, if resolutions provide reasons we can apparently give ourselves a reason to do these things merely by resolving to do so. Yet surely such bootstrapping of reasons is implausible. So, the bootstrapping argument concludes, resolutions cannot be reasons.

Footnote 11 continued

Response: there's no need for them to do so; considerations of success can support multiple views. That is, rationality could be relatively pluralistic: a disjunction of diverse ways of being successful may all be sufficient for achieving rationality. Alternatively, perhaps rationality requires different means in different circumstances. Perhaps resolutions provide reasons only to agents who cannot head off reconsideration (just as eating sugar might provide a diabetic with a reason to inject insulin without providing the same reason to a non-diabetic). Reflection on success then reveals that resolutions provide reasons for some agents.

[12] For related arguments, see Velleman (1997), Korsgaard (1997, pp. 42–50) and Broome (2001, pp. 103–105).

[13] These objections often arise in discussions of intentions (rather than resolutions) but can be applied in the current context.

We agree that such bootstrapping is implausible; we disagree that this provides grounds to abandon a (sophisticated) reasons provision account.[14] After all, just as it is typically irrational to stab one's self, it is typically irrational to resolve to do so. Consequently, in typical cases, a resolution to act irrationally could not be rationally formed.[15] Now, as RPR entails only that rationally-formed resolutions provide reasons, it does not entail that resolutions in cases like this provide reasons. So the irrational-bootstrapping argument fails.

4.2 Equal bootstrapping

Another bootstrapping argument appeals to cases involving equally-desirable options (cf. Broome 2001, pp. 98, 113). For example, perhaps we must choose between two equally-good jobs. Here, it may be rational to resolve to take one of the jobs. For example, perhaps we can't accept either job until Monday and, absent a resolution, our weekend will be ruined by flip-flopping. So we resolve to take a particular job. Now, RPR entails that we have greater reason to take that job than the other.

Why think this a problem? One might claim that it is inherently implausible that we now have a reason to choose one (equally-good) job over another. However, we simply deny this claim's force. In cases where two options are genuinely tied in all other respects, why shouldn't a rationally-formed resolution tilt the scales?

Still, the challenge can be deepened by imagining that, after we resolve to choose job 1, job 2's wage is increased slightly. Given that the jobs were previously equally desirable, job 2 is presumably now better than job 1 and so, plausibly, we should choose job 2. However, if our resolution provides us with a strong enough reason then RPR problematically appears to entail that we must choose job 1.

Still, return to the question of when it's rational to reconsider a resolution. A paradigm case where we ought to do so is when we gain new information. Further, how difficult it is to justify abandoning a resolution plausibly depends on the sort of resolution (cf. Holton 2009, p. 161). For example, a resolution to give up an addictive drug will need to be extremely robust. On the other hand, it will take little to justify abandoning a resolution to choose one of two equally-desirable options.

So when we learn about the wage hike for job 2, we gain a clear reason to reconsider our resolution. Little is required to make it rational to abandon this resolution. So we now ought to reconsider, and abandon, our resolution and choose job 2. Consequently, RPR does not entail that we should choose job 1 here and so the equal-bootstrapping argument fails.

---

[14] For prior discussion, see Hinchman (2003, p. 37) and Ferrero (2010, pp. 15–16).

[15] See footnote 3 for a discussion of the role that this sort of claim plays in our paper. In line with that footnote, we take it to be independently plausible that we typically should not form resolutions to act in a way that we now judge to be irrational. After all, the central purpose of resolutions is to ensure that our future actions are in accordance with, rather than opposed to, our current judgements about what is rational. Note that we do not make it a condition on the rational formation or retention of resolutions that they must be subjectively rational from the agent's perspective at the time that they will act in accordance with the resolution (rather than the time of resolution formation). A key role of resolutions is that they allow us to override a later perspective that we anticipate being problematic. To play this role, resolutions should be allowed to clash with other aspects of the agent's perspective at this later time.

4.3 Toxin bootstrapping

A final bootstrapping argument appeals to Kavka's (1983) *Toxin Puzzle* (see Broome 2001, pp. 101–102). In this puzzle, a millionaire offers you $100,000 if you resolve now to drink an unpleasant, but harmless, toxin tomorrow.[16] Here, it is plausibly rational to form the resolution, because doing so ensures that you receive $100,000. However, it is plausibly irrational to drink the toxin the next day, because by then you either already have the money or don't. Either way, why drink an unpleasant toxin?

Now many think that the toxin resolution cannot rationally be formed. If so, RPR will not entail that you have a reason to drink the toxin and so will not conflict with the claim that it is irrational to do so. Plausibly, then, we've already said enough to address the final bootstrapping argument.

Still, it's worth exploring what would follow if this view were wrong and the toxin resolution could be rationally formed. Under such circumstances, it might seem that RPR entails that you have a reason to drink the toxin. However, if this reason is strong enough it follows that you ought to do so. Insofar as this claim is implausible, we have a bootstrapping objection.

However, this is too fast: even if this resolution can be rationally formed, there are strong grounds to doubt that it may be rationally retained. To see this, note that resolutions are typically formed with an aim in mind. We might resolve to mark essays until midnight *in order to finish our marking on time*. We might resolve to drink the toxin *to get the $100,000*. So typically, resolutions have the form: "I resolve to ϕ for the sake of ψ".

Given this, we have strong grounds to abandon a resolution if we come to think that its purpose has already been achieved.[17] So if I resolve to mark until midnight for the sake of meeting a deadline (I don't believe I'll finish by midnight, but I at least need to work this much in order to meet the deadline later next day), I should abandon this resolution if I finish marking all of the essays at 11, precisely because I've already achieved the aim of meeting my deadline.

Now consider the Toxin Puzzle. The following day you already either have or don't have the money and so retaining the resolution can play no role in achieving the aim that motivated the resolution. So you ought to abandon this resolution. Then, as this resolution is not rationally retained, RPR does not entail that you have a reason to drink the toxin. The third bootstrapping objection fails. RPR survives these objections.

---

[16] Kavka's original discussion was framed in terms of intentions rather than resolutions.

[17] See footnote 3 for a discussion of the role that this sort of claim plays in our paper. In line with that footnote, we take it to be independently plausible that we should abandon resolutions if we come to believe that they've already achieved their aim (at least, under circumstances where we did not form the resolution anticipating that we would falsely come to believe that the aim had been achieved). Some will reject this claim despite its plausibility (cf. Gauthier 1994). Still, we never expected our arguments to convince everyone, regardless of their background views. We rest content with speaking to those who do accept this claim about rational retention of resolutions.

5 Arguments from imperiousness

Still, another objection can be raised against RPR: the argument from imperiousness. According to this argument, RPR fails to recognise the sense in which resolutions are overriding. Now there are various versions of this argument.[18] However, the most compelling is also one of the simplest. Consider the Cookie Monster's resolution. This doesn't seem to provide just one reason amongst many: it doesn't seem like the Cookie Monster should weigh all of the reasons for or against eating cookies, including the resolution-provided reason. Rather, the existence of the resolution seems to make irrelevant the deliciousness-based reasons for eating the cookies. Resolutions seem to override other reasons, rather than simply adding additional reasons into the mix. Yet RPR seems ill-placed to capture this fact.

5.1 Resolving the argument

However, further reflection reveals that imperiousness concerns do not undermine RPR but instead tell us something about the sorts of reasons that resolutions provide. In particular, such concerns push us towards a version of RPR that has previously been discussed by Raz (1975b, pp. 65–71; 1975c, pp. 489–494) and Ferrero (2010).[19]

To outline this, we need to distinguish between first and second order reasons (see Raz 1975a, b, c). First-order reasons are the standard reasons that we tend to discuss. Bad health effects, for example, might provide a first-order reason to refrain from eating arsenic. Meanwhile, second-order reasons operate by changing the normative relevance of first-order reasons. So such a reason might make normatively irrelevant some first-order reason, or might make some first-order reason more or less weighty.

As Raz (1975a, p. 164) notes, one might think of the reasons involved in consenting to harm in this second-order way. That is, if a friend consents to me causing her minor pain (perhaps as in various television prank shows), the fact that it will hurt her still seems to provide a reason against the action. However, my friend's consent might provide a second-order reason to disregard the minor-pain-related reasons against the action.[20] Now, second-order reasons are in some sense overriding: they aren't just extra reasons to be weighed alongside others but rather determine the relevance of other reasons. So if resolutions provide second-order reasons then the problem of imperiousness is resolved: second-order reasons aren't just one more reason thrown into the mix but instead work by overriding other reasons.

---

[18] Bratman (1987, pp. 23–27) suggests that RPR fails to capture how resolutions lead agents to take certain ends as fixed in further reasoning. However, we see little reason to accept this (particularly as RPR need not provide the full story of resolutions). So we focus on what we take to be the more serious challenge. Note that the following response also applies to a variant argument in Broome (2001, pp. 107, 113–114).

[19] Also Hinchman (2003). While we don't fully endorse these proposals, our account is in the same broad spirit as these prior accounts. As such, we are not presenting a novel account but rather arguing for an unduly-neglected view. Note that Bratman (1987, p. 178) claims that such views fall to bootstrapping problems, so we take the above discussion to advance the case for such views by showing that Bratman is wrong here.

[20] Whether or not it actually does so is not of great concern to us (as we use the example purely illustratively). For other examples, see Raz's works cited above.

This point can be taken more slowly. On the picture of rationality we'll assume in this paper, a decision is rational if it's supported by the weight of reasons. So we have a weighing function that takes as input a set of reasons and that outputs a weight-of-reasons-support for each possible decision. First-order reasons get their normative force from the fact that, under the right circumstances, they are fed into this weighing function. The problem of imperiousness results from the fact that resolutions don't function simply by feeding something into this weighing process in the same way; instead, resolutions seem to play some sort of more-overarching role. So resolutions cannot operate by providing first-order reasons, as these gain their normative force in the wrong way.

However, second-order reasons do not get their normative force in the same way as first-order reasons; they are not part of the input to the weighing function. Instead, second-order reasons gain their force by influencing the first-order reasons to be fed into the function.[21] So if resolutions provide second-order reasons then resolutions are imperious: they're not simply fed into the process like any old first-order reason but instead play a more-overarching role. Again, the problem of imperiousness has been addressed.

5.2 Making things precise

To make this account of resolutions precise, we first note that Razian second-order reasons are reasons to do (or not do) something *for some particular reason*.[22] For example, a second-order reason might provide a reason not to consume cookies for the reason of their deliciousness. Now, in order to be effective, resolutions will often need to provide multiple second-order reasons. After all, there might be multiple first-order reasons to consume the cookie: both the taste and the texture, say. So resolving to avoid cookies will need to provide second-order reasons not to consume the cookie for either of these reasons. With this in mind, we are led to a Razian precisification of RPR:

> RPR*: A rationally-formed, rationally-retained resolution to ϕ provides a set of second-order reasons that jointly reweight the first-order reasons so that these now support ϕ-ing.[23]

Note that this isn't a replacement for RPR, which was neutral on what sort of reasons resolutions provide; RPR* doesn't replace RPR, but simply makes it more precise. To

---

[21] In an extreme case, a second-order reason might stop a first-order reason from being fed into the function at all. In other cases, a second-order reason might allow a first-order reason to play a role in the weighing process, but might weaken (or indeed strengthen) the weight it's to be given in this process. For a reason to operate in this way is simply what it is to be a second-order reason, so this claim is true as an analytic matter. Of course, one could deny that such reasons exist, so we rely on Raz's arguments that they do (Raz 1975b). For ourselves, we find these arguments compelling.

[22] Though see Moore (1989) for a discussion of ambiguities in how Raz presents this notion.

[23] The resolution doesn't necessarily reweight all first-order reasons. Rather, resolutions will target some set of anticipated, future first-order reasons. Sometimes the targetted set might be quite narrow. For example, the Cookie Monster might target just the reasons related to deliciousness of cookies. Sometimes the targetted set might be quite broad (perhaps even as broad as all possible reasons that could support some action).

put this another way, RPR* can be thought of as a conjunctive claim, with RPR as the first conjunct and the second conjunct stating that the reasons provided by resolutions are second-order in nature.

RPR, the first conjunct, continues to be supported by the pragmatic and platitude arguments, which provided grounds to think that resolutions provide reasons without favouring any view about the nature of these reasons. The pragmatic argument relied on the fact that it's often reasonable to conclude that someone's success reflects their responsiveness to reasons. This success could just as well reflect a responsiveness to second-order reasons as first-order ones. Then the platitude argument relied on an appeal to reasons being able to make sense of the platitude that it's often rational to act in accordance with our resolutions. A second-order reasons account can make sense of this. After all, on a plausible version of such a view, the second-order reasons provided by a resolution won't just be random but will be targeted: they will reconfigure the first-order reasons in such a way that these now support the resolved-upon action. Such a view can account for the platitude.

So much for the first conjunct. As outlined above, the second conjunct is then supported by its capacity to resolve the argument from imperiousness.

Further, our defence against bootstrapping arguments applies just as straightforwardly to RPR* as it did RPR. RPR* still takes resolutions to provide reasons only if they are rationally formed and retained. Insofar as it was this feature that allowed RPR to avoid the bootstrapping challenges, RPR* also avoids these challenges. For example, in responding to the irrational form of the bootstrapping argument, we noted that it's typically not rational to form a resolution to do a clearly-irrational thing (like stabbing ourselves). As only rationally-formed resolutions provide reasons, a resolution to stab ourselves provides no reasons whatsoever and so there is no bootstrapping of reasons here. In considering this response, it hardly matters what sort of reasons we're focusing on: if no reasons whatsoever are provided by this resolution then it's irrelevant what sort of reasons would be provided by a rationally-formed resolution.[24]

In the light of these considerations, we adopt RPR*.

6 Conclusions

The idea that resolutions provide reasons is a natural one. It accounts for the normative force of resolutions and accords with the way we speak when we justify acting in

---

[24] Indeed, an additional response can be given to provide RPR* with an even stronger defence against the bootstrapping objections. Note two things. First, resolutions don't change the options available to the agent but instead change the agent themselves. Second, first-order reasons seem to capture the considerations that favour or disfavour options. So if resolutions provide first-order reasons then we have a case where the options remain the same but the considerations favouring them change. This might seem odd, and this oddity might be taken to give the bootstrapping objection much of its force. On the other hand, second-order reasons don't capture considerations that favour or disfavour options, but rather tell us which of these considerations is relevant to the agent. So if resolutions provide second-order reasons then we have a case where: (a) the options are the same and so too are the considerations favouring the options; but (b) the agent is different and so too are the considerations that are relevant to the agent's normative perspective. This might seem much less odd, and so this underlying source of force for the bootstrapping objections falls away. Not only do our initial responses to the bootstrapping objections remain, but we now have additional grounds to be unconcerned.

accordance with resolutions. However, there is something of a consensus that the reasons provision view is fatally undermined by bootstrapping arguments and the argument from imperiousness. This is where we enter the debate. In particular, we have argued that sophisticated versions of the reasons provision account can avoid both objections. Given that there are also positive arguments for accepting such a view, we conclude that RPR* has been unfairly denigrated. Resolutions provide (second-order) reasons. The Cookie Monster can rationally quit cookies.

References

Bratman, M. (1987). *Intention, plans, and practical reason*. Cambridge: Harvard University Press.

Broome, J. (2001). Are intentions reasons? And how should we cope with incommensurable values. In W. M. Christopher & A. Ripstein (Eds.), *Practical rationality and preference: Essays for David Gauthier* (pp. 98–120). Cambridge: Cambridge University Press.

Cohen, D., & Handfield, T. (2010). Rational capacities, resolve, and weakness of will. *Mind*, *119*, 907–932.

Ferrero, L. (2010). Decisions, diachronic autonomy, and the division of deliberative labor. *Philosophers' Imprint*, *10*, 1–23.

Gauthier, D. (1994). Assure and threaten. *Ethics*, *104*, 690–721.

Hinchman, E. S. (2003). Trust and diachronic agency. *Noûs*, *37*, 25–51.

Holton, R. (2009). *Willing, wanting, waiting*. Oxford: Oxford University Press.

Kavka, G. S. (1983). The toxin puzzle. *Analysis*, *43*, 33–36.

Korsgaard, C. M. (1997). The normativity of instrumental reason. In G. Cullity & B. Gaut (Eds.), *Ethics and practical reason*. Oxford: Clarendon Press.

Moore, M. S. (1989). Authority, law, and Razian reasons. *South California Law Review*, *62*, 827–896.

Quinn, W. S. (1990). The puzzle of the self-torturer. *Philosophical Studies*, *59*, 79–90.

Raz, J. (1975a). Permissions and supererogation. *American Philosophical Quarterly*, *12*, 161–168.

Raz, J. (1975b). *Practical reason and norms*. London: Hutchinson.

Raz, J. (1975c). Reasons for action, decisions and norms. *Mind*, *84*, 481–499.

Velleman, J. D. (1997). Deciding how to decide. In G. Cullity & B. N. Gaut (Eds.), *Ethics and practical reason* (pp. 29–52). Oxford: Oxford University Press.

Verbeek, B. (2014). On the normativity of intentions. *Topoi*, *33*, 87–101.