



ARTICLE

Normative Resilience

Henrik Andersson*  and Jakob Werkmäster* 

Lund University, Lund, Sweden

Emails: Henrik.Andersson@fil.lu.se, Jakob.Werkmaster@fil.lu.se

Abstract

This article discusses the phenomenon of normative resilience, with a focus on evaluative resilience. An object can become more or less valuable. In addition to this change in an object's value, the object's value can become more or less resilient. If it is less resilient, it cannot withstand as much evaluative change without its degree of value changing, as compared to an object with more resilient value. The article consists of three parts. First, examples of resilience are presented to give the reader an intuitive understanding of the phenomenon, Second, the Fitting Attitudes Analysis of value is invoked to provide a formal account of evaluative resilience. Third, the theoretical and practical advantages of acknowledging the existence of evaluative resilience are brought to light.

Introduction

The aim of this article is to highlight a feature of the normative landscape, that goodness and deontic properties, such as 'ought' and 'permissibility', can be more or less *resilient*. Resilience accounts for a dimension of the normative that has previously been overlooked; not only can an object be more or less good, but its degree of value can also be more or less resilient, and not only can an act be either permissible or impermissible, but its permissibility or impermissibility can also be more or less resilient. The notion of resilience is easy to grasp in the deontic domain. That certain conditions can solidify the permissibility of an act should not be surprising. The idea of evaluative resilience, however, may be more thought-provoking. Normally, we think that an evaluative change entails that the value of an object has increased or decreased. In this article, we show that there are cases where a change to an object's value need not make the object better or worse. This new dimension of value is philosophically sound and can help shed light on interesting topics within value theory, by offering a new way of explaining notions such as 'evaluative tolerance' and 'clumpy values', which in turn can bring insights to arguments that involve small evaluative changes of any kind.

The structure of the article is as follows. In the first part of the article, we provide examples that establish that normative resilience exists. In the second part of the article, we show that an influential analysis of value already provides the conceptual space

*Henrik Andersson and Jakob Werkmäster contributed equally to this work. Their names are listed alphabetically.

© The Author(s), 2022. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

necessary to model the resilience of value. In the third part of the article, we show how resilience forces us to reconsider some crucial notions in the normative discourse, and how the concept sheds new light on contemporary discussions in value theory. The article concludes by arguing that the arguments presented show that normative resilience is an intelligible and useful feature of the normative landscape.

Part 1: The case for resilience

In the introduction, we presented a rough characterization of what we take to be resilience; an object's value is resilient if and only if it can withstand an evaluative change without becoming more or less valuable. We believe that situations of this kind are familiar to the reader. That is, a change to the properties that the value of the object depends upon has taken place and yet we hesitate to say that the object has become better or worse. The resilience of value can be contrasted with physical resilience; an object can be physically resilient if it can withstand physical force. A rock is in this sense more resilient than a glass vase. Similarly, an object's value is more resilient if it can withstand more evaluative change. This means that the resilience of value need not be linked to the physical resilience of the bearer of the value. For example, Tommy is indeed valuable due to the fact that he is a human being. Even if Tommy were to accidentally saw off his index finger, his value as a human being would remain completely unchanged. This is because not all physical changes of a value bearer entail an evaluative change. In this article, we are solely concerned with evaluative changes.

An object can undergo an evaluative change in many ways. With the change of an object's natural properties, or a change to natural extrinsic properties of the object, many evaluative changes are possible. For example, the valence of the object's value can change by transforming a good object into a bad object; the degree of value can change, a valuable object can become even better (or less good). Further, the type of value an object has can change from being, say, contributory valuable to being finally valuable, and the kind of value an object has can change from being, say, an aesthetic value to a moral value. This list is, however, not exhaustive. In this article, we propose that normative resilience is a type of evaluative change that has previously been neglected in the literature.

It must also be clarified that the resilience we have in mind is primarily to be construed as a dimension within a specific category of value. For example, while it may be possible for an object to undergo a change so that it is transformed from having solely a contributory value to having a final value, this is not the change we primarily have in mind. We have chosen to focus on the resilience within each kind of value. That is, an object can, for example, have a certain amount of final value and a certain amount of contributory value, and the object's final value can be more or less resilient while the object's contributory value need not be resilient to the same degree.¹

It is hard to deny that there can be deontic resilience. For this reason, we will start by considering deontic resilience. When the existence of deontic resilience has been established, we will shift our focus to the more interesting idea of evaluative resilience.

Resilience is a property that is frequently indirectly referred to when it comes to moral oughts or epistemic oughts. Simply put, sometimes what you ought to do is quite straightforward: one, or several, of the reasons in virtue of which you ought to

¹We are grateful to an anonymous reviewer for presenting this possibility to us. While what we say here is applicable to most kinds of value, we, for brevity and clarity, use the notions such as 'value' and 'good' to refer to final value and what is finally good.

do something could be absent, and it would nonetheless be the case that you ought to do it. This can be illustrated by the following example.

Resilient oughts

If you press a button, you will receive a delicious cake and a substantial sum of money. Since pressing the button comes at no cost to you or anyone else, your reasons against pushing the button are weak or even non-existent. The fact that you will receive a delicious cake is a strong reason to press the button. The fact that you will receive a substantial sum of money is also a strong reason to press the button. You correctly deduce that you ought to press the button.² In this scenario, the ought is quite resilient. That is, even if the circumstances change so that there now is a weak reason against pushing the button (perhaps you will feel slightly nauseous of the sugar rush that the cake will give you), it would still be the case that you ought to press the button. Similarly, even if we removed or weakened one of the reasons to press the button, say by removing or lowering the monetary reward, you still ought to press the button.

The fact that you will receive both delicious cake and a substantial sum of money does not mean that you ‘more ought to’ press the button than if you simply received the delicious cake. Ought is not gradable in this sense. You simply ought to press the button. Still, the deontic situation is different in the case where you only receive delicious cake and the situation where you receive delicious cake and a substantial sum of money. What is different, we argue, is how resilient the ought is.

There are similar examples for the concept of ‘permissibility’. An act cannot be more or less permissible: it is either permissible or not. Yet, for reasons similar to those presented when discussing oughts, permissibility can in one sense be understood to come in degrees since it can be more or less resilient. An act whose permissibility is very resilient is in a sense further away from being impermissible.

Similarly, we encounter resilience when we are dealing with epistemic oughts. It is not unusual that the evidence overdetermines the justified belief. Say that a judge has conclusive reason to believe that it is beyond reasonable doubt that Eija is guilty of the crime in question. The judge ought to have this belief and come to this verdict. The judge can hold the belief with full credence, that is, have an outright belief that Eija is guilty. Now, assume that the technical analysis of the crime scene is admitted to the court as evidence. The technical analysis also points to the fact that Eija is guilty. Even though the judge already ought to believe that Eija is guilty, the addition of the technical analysis makes this epistemic ought even more resilient. It cannot be that the judge ought to hold the belief with a higher credence, since it was already the case that the judge should have a full belief that Eija is guilty. That she ought to hold the belief with full credence is, rather, more resilient – resilient in the sense that more evidence would have to be presented to make it the case that the judge no longer ought to believe that Eija is guilty.³

²We here help ourselves to the common assumption that there is a connection between what you have reasons to do and what you ought to do. Most commonly, the connection is understood as it being the case that you ought to do what you have most reason to do (Parfit 2011). For someone who denies that for every ought there must be a reason, see Horty (2012).

³One might also argue that the judge’s belief is now safer given the technical analysis. However, using a semantic of safe beliefs that allows for degrees of safety might be controversial. For more on this see Rabinowitz (2011). For more on the similarities between safe beliefs and resilience, see n. 16 below.

The dimension of resilience is not only to be expected in the deontic domain. We take the deontic domain to be a subset of the more general normative domain.⁴ This makes it plausible that resilience plays an important role in the axiological domain as well. A first step has thus been taken to establish evaluative resilience. We will now move on to consider convincing examples of evaluative resilience. These examples are similar to the deontic examples in the sense that the values considered do not come in degrees but are either actualized or not. The notion of resilience allows us to acknowledge that there is another dimension for which these values can fluctuate without necessarily affecting the fact that they are actualized. Which values are binary in this sense is up for debate in the field of normative ethics. Nonetheless, some plausible examples can be given.

Human dignity

Every human has the value of dignity that should be respected. This value is different from the value other beings have in virtue of being able to experience pleasure and suffering. Humans have this value in virtue of our rational capacities and our possibility to lead autonomous lives. Rational capacities are relevant to our human dignity. Bill and Valter differ in their rational capacities but when it comes to their human dignity, they are equally as valuable. The resilience of their dignity can thus depend on their rational capacity and could consequently differ between them and yet we can agree that they both have the value of dignity to the same degree.

It is of course debatable what level of rational capacity is required in order to have human dignity. Regardless of what that level is, we can still imagine someone's rational capacity deteriorating in such a way that it comes closer to the point where it would be, at least, indeterminate whether she retains her human dignity. In the example above, Bill and Valter have the same human dignity, but a greater deterioration of, say, Bill's rational capacity would need to take place for Bill to lose said value than in the case of Valter.

No. 1

Alexander and Gustaf compete in two 100-meter races. In the first race, Alexander wins with an excellent time. Alexander finishes first in the second race as well, albeit this time with a slower time. After the race, Gustaf comments: 'I almost got you this time, you're not that much better a runner than me'. Alexander replies, 'A win is a win. How quickly we run relative to each other is relevant to who's number one, but as far as I am concerned, what matters is who won. And my win in this second race is no less a win than my win in the last race'. The artificial value of 'being No. 1 in a race' accrues to Alexander in both races. In the second race, however, it is less resilient.

The example of No. 1 can easily be reiterated following a similar structure. A coarse-grained value is determined by a more fine-grained criterion. In the case of the races, we achieve an ordinal scale (placement in the race) based on something that is given in a cardinal scale (run-times). It might be that according to some more fine-grained value, Alexander is less good in the second race, but given the assumed comparison scale, this is irrelevant.

⁴Here we follow the common view that in the domain of normative concepts two subsets of concepts can be differentiated: evaluative concepts (e.g., good, bad, better) and deontic concepts (e.g., ought, permissible, forbidden), see Tappolet (2013). We are grateful to an anonymous reviewer for suggesting this point of clarification. For another way to divide the normative domain in terms of what there is reason to do – that is, feel, believe, and act – see Skorupski (2010).

These examples could be accused of being trivial and it could be claimed that there are other ways to account for the phenomenon that the examples highlight. There are, however, much-discussed cases in value theory where the concept of resilience plays an important role and where the accusation of being trivial does not stick. For example, cases that are sometimes explained by referring to the tolerance of the equally as good as relation. The tolerance property of ‘equally as good’ amounts to the fact that for two objects that are equally as good the objects tolerate some evaluative change without causing one of them to be better than the other. That ‘equally as good’ is tolerant can be corroborated by referring to our everyday use of the term. When we use the term, we do not necessarily mean that the things we are comparing are precisely identical in value. We can admit that there is an evaluative difference between objects but this does not amount to the fact that one is better than the other one is. While it has previously not been possible to explain what makes a relation tolerant, the notion of resilience can now fill that explanatory gap. The equally as good as relation is often very resilient and thus it can tolerate evaluative changes. Rather than having to postulate that the equally as good relation differs significantly from other value relations in this respect, we are able to give a better and more general explanation of tolerance by appealing to evaluative resilience.

An idea somewhat similar to the notion of tolerance is the suggestion that comparisons can be ‘clumpy’, which means that it may very well be that two objects are not exactly equally as good but for the comparison at hand, the two objects are sufficiently similar for it to be correct to say that they are equally as good (Hsieh, 2005; Peterson, 2007). That is, finegrainedness is not always needed and thus we can sort elements in a sequence into clumps and treat elements in this clump as equally as good. Clumpiness was first introduced as a means to account for so-called hard cases. More specifically, Nien-hê Hsieh (2005) argues that because there is no need to assume that there are more value comparatives than the ‘at least as good as’ comparative to account for hard cases, the clumpiness of value is sufficient. We find the notion of clumpiness to be interesting since it, from an explanatory point of view, is similar to the concept of ‘resilience’. Adherents of the clumpiness view believe that objects in the clump differ evaluatively from each other and yet they can be treated as equally as good, so far as the views correlate. However, we suggest that the objects not only can be treated as equally as good but also in fact *are* equally as good. Their evaluative difference ought instead to be accounted for in terms of resilience.⁵

We will have reasons to return to the discussion on evaluative resilience and value comparisons later on when we show that resilience can play a helpful role in shedding light on intricate discussions within value theory such as the Small Improvement argument and various spectrum arguments. First, however, we provide a formal account of the resilience of value.

Part 2: The Fitting Attitudes analysis of value

So far, we have presented examples of resilience in order to provide the reader with an intuitive understanding of the concept. The groundwork has, however, not yet been done. That is, we have not shown that there is conceptual space for resilience yet. In this part of the article, we adopt the influential Fitting Attitudes Analysis of Value (henceforth the FA analysis) to show that there is such conceptual space and also provide an account for the resilience of the value. By showing that there is conceptual space

⁵The insightful comments from an anonymous reviewer led us to revise our previous view on the relation between clumpy values and evaluative resilience. We are very thankful for this feedback.

for this feature of a value and providing an analysis of resilience, we improve the plausibility of the idea that there is such a thing as normative resilience.

According to the highly influential Fitting Attitudes analysis of value, an object being valuable is understood in terms of it being fitting to adopt pro-attitudes or con-attitudes towards the object. If it is fitting to adopt a pro-attitude then the object has a positive value and if it is fitting to adopt a con-attitude then the object has a negative value.⁶ This entails the following biconditional.⁷

*FA analysis: x is valuable if and only if it is fitting to favor/disfavor x.*⁸

The analysis has two components, a normative component, and an attitudinal component. In the definition above the normative role is played by ‘fitting’ but it is also often spelled out in terms of ‘reasons’ or ‘ought’ (Ewing 1948; Rabinowicz and Rønnow-Rasmussen 2004; Hieronymi 2005; Rabinowicz 2008; Schroeder 2010; Heuer 2011). For simplicity, we will follow suit and formulate the normative component in terms of reasons. The attitudinal component ‘favor’ or ‘disfavor’ is often used as a placeholder for any type of pro-attitude (or in the case of negative value, a con-attitude) such as admiring, loving, appreciating, and so on.

The FA analysis has primarily been construed as an account of what it is for something to be valuable, but for the account to be a full-fledged analysis, it must also be able to account for the phenomenon that value comes in degrees.⁹ There seem to be two obvious ways for the adherents of the FA analysis to develop their account so that it can encompass this feature of value. A first possibility is to describe degrees of value by the strength of the normative component and the second possibility is to describe it by referring to the intensity of the attitude. According to the first line of analysis, if *x* is very valuable there are *very strong reasons* to favor *x*, while according to the second proposal there are reasons to *favor x a lot*. There seems to be some intuitive support for both accounts.¹⁰ Intuitively, if we have very strong reasons to favor something

⁶Thus, our use of the term ‘value’ is neutral between whether the value in question is a positive value or a negative value.

⁷The FA analysis of value is one of the more dominant theories of value, popularized by Scanlon (1998) but goes back to Brentano (1969 [1889]) and Ewing (1948). For more developments that are contemporary, see among others Rabinowicz and Rønnow-Rasmussen (2004), Rønnow-Rasmussen (2011), and Rowland (2019).

⁸We use ‘favor’ as a generic term for a pro-attitude. The FA analysis of value is a hotly debated theory and faces several prominent objections, such as the Wrong-Kind of Reasons problem (for a recent overview see Gertken and Kiesewetter (2017), and the solitary goods objections (Bykvist 2009) to name just two. We believe that nothing in this article is committed to a particular stance concerning these other issues. For this reason, these objections will not be addressed in this article.

⁹For more on this see Andersson and Werkmäster (2021).

¹⁰A third possibility is that both components together determine degrees of value. As far as we are aware, no one has opted for this interpretation of the FA. Such an interpretation requires an explanation of how the two components interact. Such an explanation seems hard to come by, which might be why no one in the literature has opted for this possibility. As was pointed out by a reviewer, there is also precedence in the literature on the FA analysis (Brentano 1969 [1889]; Rabinowicz, 2008) to understand the betterness relation in terms of what it is fitting to prefer, that is, *x* is better than *y* if and only if it is fitting to prefer *x* to *y*. We agree with such an analysis of comparative value. Such an analysis, however, is an analysis of comparative value and not straightforwardly applicable to monadic value. In this article, we are primarily interested in monadic value. In order for the analysis of comparative value to be applicable in the present article, one would have to assume that comparative ‘better-than’ can account for the monadic ‘good’. It must also be pointed out that, understanding *x* being better than *y* in terms of it being fitting to prefer *x* over *y* is

then this should, somehow, be reflected in the value of the object.¹¹ On the other hand, if there are reasons to favor something to a high degree then this should also manifest itself in the value of the object.¹²

Both accounts open up the conceptual space necessary for modeling the resilience of value. If the normative component accounts for the degree of value, then the strength of the attitudinal component can account for degree of resilience. If the attitudinal component accounts for degrees of value, then the strength of the normative component can account for the degree of resilience.

Andersson and Werkmäster (2021) and Rabinowicz (2020) have recently advanced support for the approach that emphasizes the attitudinal component. For ease of presentation, we will assume that the attitudinal component of the FA analysis accounts for degrees of value. Consequently, we will utilize the normative component to model the concept of the resilience of value.¹³

Resilience according to FA

That there are two dimensions of value according to the FA is best illustrated by an example. Consider two objects, x and y . There is a conclusive strong reason to favor x with intensity S and a conclusive weak reason to favor y with intensity S . This is illustrated in Figure 1.

For both objects, the reasons to favor the objects are stronger than those against favoring are. That is, both objects are good.¹⁴ Furthermore, since the degree of value

compatible with both of our suggestions in the main text. At least it is undeniable that we can have stronger or weaker reasons to favor something, and that we can favor something to a higher or lower degree of intensity. All that we need in order to give an analysis of resilience is that these two parameters exist. It is beside the point if it turns out (or not) that it is, say, fitting to favor x to a high degree of intensity *because* it is fitting to prefer x over y or vice versa.

¹¹More precisely, if we have very strong reasons *of the right kind* to favour something, then this should (somehow) be reflected in the value of the object. Enough ink has been spilled over the Wrong Kind of Reasons (WKR) problem for the FA analysis. In this article, we have chosen to set this problem to one side. For the rest of the article, we will assume that the reasons mentioned are of the right kind.

¹²It should be noted that some theorists accept a view that seems to entail that the normative component does not come in degrees. They treat the normative component as an ungradable concept. Prime examples may be those who refer to the normative concept as ‘ought’, ‘sufficient reason’, or ‘conclusive reason’ (Brentano 1969 [1889]; Rabinowicz 2008; Skorupski 2010; Gertken and Kiesewetter 2017). To have a conclusive reason to favor an object means that the reasons for having the attitude outweigh the reasons against having the attitude. From this it follows that one cannot have ‘more conclusive reasons’ or ‘less conclusive reasons’. The concept does not come in degrees: either the reasons are conclusive or they are not and the same goes for ‘ought’ and ‘sufficient reason’. Likewise the predicate ‘wins the race’ does not come in degrees in *No. 1*. If Alexander finishes first in one race, he wins that race. If he then goes on to finish first in another race, but with a better time than the previous race, he does not win the latter race *more* than the first one. He simply wins both races. However, what is important here is that there is still a difference in Alexander’s speed in both races and how fast he ran, just as there can be a difference in how much the reasons to favor an object outweigh the reasons against favoring the object. That is, even with this formulation of the FA it seems to be possible to account for the degree of value by referring to the strength of the normative component.

¹³Note, however, that the same result (i.e., that the FA analysis provides the conceptual framework necessary for giving a characterization of resilience) is possible by modelling degrees of value via the normative component and the resilience of value via the attitudinal component, *mutatis mutandis*.

¹⁴To make it explicit, we use ‘favor’ as a placeholder for a positive pro-attitude and ‘disfavor’ as a placeholder for a negative attitude. Given that one cannot hold a positive attitude with a negative intensity, any object that it is fitting to favor is good. Any object that it is fitting to disfavor is bad. This being said, it is of

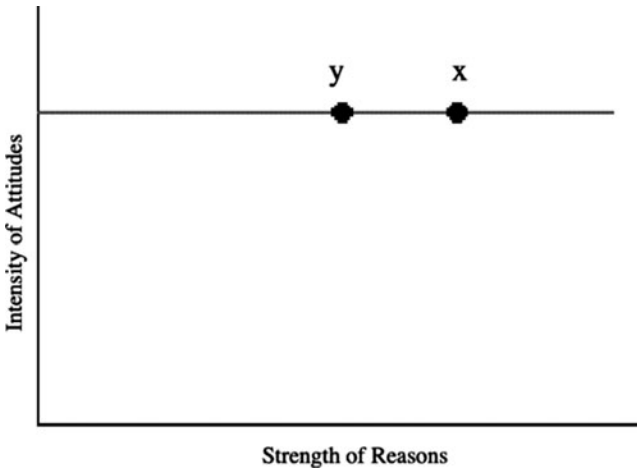


Figure 1. Intensity of attitudes and strength of reasons.

is determined by the strength of the attitude it is clear that both x and y are valuable to degree S . Simply put, they are equally as good. One might find this odd since the normative component differs between x and y . For x the normative strength seems robust, that is, a drastic change to the balance of reasons must take place for the overall balance of reason to change while the same cannot be said for y .¹⁵

Figure 2 illustrates how this amounts to a dimension of value. At the left-hand side of the horizontal axis, the strength of the reason in favor of x is weaker, and at some point, this reason is outweighed by reasons that are incompatible with favoring x . On the right-hand side of the horizontal axis, the reason to favor x is very strong, making x 's value V very resilient.¹⁶




Figure 2. Non-valuable and valuable objects.

course nonetheless fitting to prefer a lesser evil to a great evil. It being fitting to prefer a lesser evil to a greater evil does not entail that it is fitting to favor, that is, hold a pro-attitude towards, the lesser evil *tout court*. Thanks are owed to a reviewer for pressing us on this issue.

¹⁵It is possible to argue that this should not be spelled out in terms of resilience but rather that these are different values. For each *pro tanto* reason to have an attitude towards an object, the object will have a corresponding value. This seems, however, highly unlikely. With such an account, there would be too many values to ascribe to the object. In order to avoid this consequence, proponents of the FA analysis should be motivated to accept our view or abandon the FA analysis.

¹⁶As has been suggested by a reviewer, one could perhaps question the relationship between the strength of reasons and resilience. Resilience is a counterfactual and thereby modal notion. To illustrate why the resilience of an object's value (or in the following case an ought) might not correlate with strength of reasons consider the following. You ought to take out the trash. You ought to do this in virtue of the very

What we can tell so far is that the resilience of value is not the same as the degree of value. However, what else can be said of this feature of value? Luckily, the FA account gives us some clues on how to understand this feature. According to the account, an object, x , is more resilient in value than an object, y , if and only if a more substantial change to the reasons for favoring x is needed than compared to y in order to change the degree of value of the objects. That is, the strength of the reasons against favoring x can increase (or the strength of the reasons for favoring x can decrease) more than in the corresponding case of y . Roughly put, the more the strength of reasons for and against favoring an object can decrease/increase without a change in value, the more resilient the value of the object is.¹⁷

The property of being resilient is thus not fully orthogonal to the property of being valuable. It is the object's value that is resilient. If the value is determined to not be resilient there are two possibilities available: (i) it could mean that the reasons to strongly favor the object just barely outweigh the reasons to favor the object even more strongly, or (ii) it could mean that the reasons to strongly favor the object just barely outweigh the reasons to favor the object less strongly (or not favor it at all). This is illustrated in Figure 3.

In the middle of the horizontal axis the balance of reasons (more or less strongly) favor favoring the object to degree A , and hence has value V . The arrows at the top illustrate the balance of reasons to favor the object with more or less intensity. On the left-hand, side the reasons to favor the object with a less strong intensity, $A-$, almost outweigh the reasons to favor the object with intensity A . Here the object is close to having a lesser value, $V-$. Similarly, on the right-hand side the reasons to favor the object with a stronger intensity, $A+$, almost outweigh the reasons to favor the object

strong reason that your boss ordered you to do so. She, however, decided that you should take out the trash rather than your other colleague by a coin toss. In which case, there is a nearby and physically similar world where you do not have this reason and therefore are not obligated to take out the trash. The fact that there is a very nearby possible world where it is not the case that you ought to take out the trash might lead one to suspect that the ought was not very resilient in the first place, and hence reject the connection between the strength of reasons and resilience. We want to resist this conclusion. While it is true that the nearby possible worlds in which it is the case that you ought to take out the trash and in which it is not the case that you ought to take out the trash are physically very similar, we want to say that the ought is nonetheless resilient. It is resilient because normatively speaking the worlds are not at all that similar. A small change in the supervenience base can entail a big normative difference. A more relevant possible world to compare it with is the normatively similar world where your boss ordered you to take out the trash but this would entail you missing the final minutes of the World Cup final. With that said, and as was demonstrated, it is possible to model resilience as a possible worlds ordering, where the resilience of an object's value or an ought is determined by the size of the set of normatively similar worlds in which the object is valuable, or in which you ought to do something. This echoes similarities to how one can determine the safety of a belief (Sosa 1999; Williamson 2000; Pritchard 2009; Whiting 2020). The major difference between safe knowledge and resilient values concerns which possible worlds are relevant when determining safety and resilience. The reason why we refrain from modelling resilience in terms of a possible worlds ordering is to avoid confusion about which nearby worlds we should use as comparisons. The possible worlds should not be ranked based on how physically similar they are but rather how normatively similar they are. We are very grateful to an anonymous reviewer for pushing us to think about these issues.

¹⁷We have not specified what kind of value we have in mind. The FA analysis can account for a plethora of values, such as goodness *simpliciter*, goodness of a kind, personal goodness, impersonal goodness, and so on. For all of these accounts of value it is also possible to account for the value's resilience. We can thus be neutral to substantive views of values since whatever view one has it is possible to account for the resilience of the relevant values.




Figure 3. Resilient to improvements or worsenings.

with intensity A . Here the object is close to being of greater value, $V+$. At the far-left side of the axis, the object's value is very resilient in the sense that the object is far from becoming better, but it is not resilient to becoming worse. The opposite will be true for the right-hand side of the illustration and in the middle; the object's value is very resilient in the sense that it is quite stable.

Part 3: Applicability

Our main aim in this article is to show the conceptual feasibility of objects' values being more or less resilient. Over and above gaining a deeper understanding of value, the notion of resilience has a practical use. Below we show that resilience is applicable and helpful in discussions of value comparisons, moral philosophy, and decision theory. In short, in decision theory, it opens up the possibility to deny the principle that an option is rationally required if and only if it is *better* than all other options. In moral philosophy, resilience allows us to explain how something can undergo an evaluative change without needing to become better or worse. This becomes especially relevant in how to understand collective harm cases and calls into question certain arguments for the possibility of a fourth *sui generis* value relation.

Proof of Concept 1: Resilience or parity in the Small Improvement Argument

By expanding our normative toolbox with the notion of 'resilience' it is possible to account for many puzzling results within normative theory. Below we show possible ways in which resilience can account for some of the intuitions that are involved in cases where one is asked to consider a small evaluative change that is being done to an object.

Consider, for example, the *Small Improvement Argument* that is often invoked in order to reach the conclusion that the standard three value relations, that is, better than, worse than, and equally as good, cannot fully account for all value relations. The argument takes roughly the following form. Imagine A and B , such that A is not better than B nor is B better than A . Furthermore, imagine a slightly improved A , $A+$. For some value bearers, it is possible that $A+$ is not better than B . If A and B are equally as good, then $A+$ must be better than B , and thus we can conclude that A and B are not equally as good (see Chang 2002; Gustafsson 2013). Examples are then given to show that there are objects that stand in the relation previously described. For example, consider the choice between two different careers, one as a corporate lawyer and one as a philosopher (Chang 2002). It is possible to imagine how some can find these careers to be so different that they judge neither to be better than the other. Furthermore, if we improve the legal career by slightly increasing the salary it could

still be that this slightly improved career is not better than the philosophical career. It then follows that the two original careers cannot be equally as good. This argument is important if one wants to argue that there are more value relations beyond 'at least as good'. Perhaps value bearers can be incomparable, or perhaps a fourth, previously overlooked value relation can be instantiated (Chang 2002). The concept of resilience allows us to question the Small Improvement Argument in a new way. For purported examples of objects that stand in the relation described by the Small Improvement Argument it is now possible to question whether the evaluative difference between A and A+ actually is to be framed in terms of 'better than'. The change that transforms A to A+ might better be captured in terms of 'resilience'. In the example with the two careers, it would be possible to argue that a slight increase in salary does not make the career overall better; it merely becomes more resilient to worsenings. The legal career with an increased salary could still be equally as good as the career in philosophy, but the resilience of their values could differ.

We are not saying that resilience can convincingly block the Small Improvement Argument. Rather we are suggesting that the concept of resilience makes the Small Improvement Argument less straightforward. Adherents of the argument can maintain that there are possible changes to the legal career that would clearly count as an improvement that makes the career overall better. But it can be objected that such an improvement may be so substantial that it would also make the improved legal career better than the philosophy career. Given the theoretical cost of having to postulate a fourth value relation, proponents of the Small Improvement Argument need to show that their small improvements actually make the legal career better rather than making its overall value more resilient to worsenings.

Proof of Concept 2: Imperceptible harms and more or less resilient harms

In cases of imperceptible harms, collective harm, and spectrum cases in general, we all share mutually incompatible intuitions. The key problem is that no individual act seems to make a normatively relevant difference, yet the collection of acts makes a large normative difference. The problem is that on each pairwise comparison the objects are equally as good (or bad). This is for example the case in Parfit's (1984: 80) harmless torturers case, in which a torturer presses a button that inflicts an imperceptible worsening of the pain suffered by a victim. However, in Parfit's example, there are thousands of torturers who press a button with the result that the victim suffers a perceptible pain. Individually, neither torturer inflicts a perceptible pain, but together they inflict a perceptible pain.

For the individual torturer, pressing the button will not make the outcome worse. Nonetheless, intuitively the torturer has a reason not to press the button. By introducing the notion of resilience, it is possible to accept that the situation is *not always* made worse by the imperceptible change and still account for how there has been an evaluative change; a change in the value's resilience. Pressing the button makes the neutral situation for the victim less resilient since it requires fewer buttons to be pushed for the situation to become bad. The outcome's value will now be less resilient to worsening.

Of course, introducing the notion of resilience does not *solve* cases of imperceptible harms but it presents a pattern of explanation that can respect the intuitions involved. Resilience allows us to do this without having to suppose either that every single

imperceptible change always makes the situation worse, or that no imperceptible change ever makes the situation worse.

Proof of Concept 3: Is it always rational to be indifferent between two equally good options?

One valuable insight is that the possibility of resilience could give us reason to question the following intuitively plausible principle:

Optimality: It is rationally permissible to choose x if and only if x is at least as good as all other alternatives.¹⁸

Assume that you are to choose between two options, x and y . Let us assume that x and y are equally as good, but the value of x is more resilient than the value of y . First, we must remind ourselves that the fact that y 's value is less resilient than x 's says nothing about whether it more easily can become better or worse. All it tells us is that y 's degree of value can withstand less normative change. Let us also assume that y is such that it is not as resilient to worsening as x is. Will this have any implications for what rationality will require us to do?

It is tempting to claim that it is preferable to opt for x . In some sense, it seems to be the safer alternative. It is, however, unclear why it is tempting and in what sense it is safer. In order for it to be preferable, x needs to have a choice-theoretic advantage over y . If it has a choice-theoretic advantage, however, it also seems natural to suppose that x is better than y , but *ex hypothesi*, x and y are equally as good and x 's value being more resilient than y 's does not affect how valuable x is. If we want to claim that x nonetheless has a choice-theoretic advantage over y , we need to deny that all choice-theoretic advantages make an option better. In other words, we would need to reject *Optimality*. Rejecting *Optimality* would be a radical position. Given the current state of evidence, it is not one we advise accepting.

The intuition seems to be that it is reasonable to opt for the alternative that is more resilient to worsenings if two alternatives are equally as good because it is in some sense safer. The question is in what sense it is safer. The intuition that it is safer, however, might simply be due to a confusion – conflating risk and expected value with resilience. One alternative's value being more resilient to worsenings does not mean that it has a higher expected value. It means that a larger normative change would need to occur in order for the alternative's value to lessen. An alternative's value being very resilient, however, might frequently correlate with the alternative being low risk, under the assumption that small normative changes are more likely than large normative changes. Given that the alternatives are stipulated to be equally as good, any risk-related considerations are already accounted for, and so treating resilience to worsenings as making the alternative safer is likely to be a case of double-counting the risk.

To give an account of the aforementioned intuition, and thereby show that it is distinct from intuitions about risk and expected value, is an intriguing challenge for normative theories on rational choice, but it is beyond the scope of this article.

¹⁸Given that rationality supervenes on the mind, one should perhaps amend *Optimality* so that it reads that it is rationally permissible to choose x if and only if x is *believed to be* at least as good as all other alternatives. For simplicity, we omit the belief condition and how such a belief condition should be spelled out. Nothing in our argument hinges on this omission.

Conclusion

This article establishes the possibility of a previously overlooked dimension of the normative: resilience. First, examples were given to establish the intuitive plausibility of normative resilience. The examples involved everyday phenomena, but examples familiar in the value theoretical literature were also discussed. Second, the article shows that the influential FA analysis provides conceptual space for accounting for what it is for an object's value to be more or less resilient. In the third part of the article, we showed the potential fields of applicability of this concept – resilience can account for certain key intuitions in many fields in philosophy. Resilience can thus not be accused of being a byproduct of a certain account of value. Rather, it is an important concept that can shed helpful light on the, sometimes mysterious, normative landscape.¹⁹

References

- Andersson, H. and Werkmäster, J. (2021). How Valuable Is It? *Journal of Value Inquiry*, 55: 525–42.
- Brentano, F. (1969 [1889]), *The Origin of Our Knowledge of Right and Wrong [Vom Ursprung sittlicher Erkenntnis]*, trans. by R. M. Chisholm and E. H. Schneewind, London and Henley: Routledge & Kegan Paul.
- Bykvist, K. (2009). No Good Fit: Why the Fitting Attitude Analysis of Value Fails. *Mind*, 118(469): 1–30.
- Chang, R. (2002). *Making Comparisons Count*. London: Routledge.
- Ewing, A. C. (1948). *The Definition of Good*. Westport, CT: Hyperion Press.
- Gertken, J. and Kiesewetter, B. (2017). The right and the wrong kind of reasons. *Philosophy Compass*, 12(5), e12412.
- Gustafsson, J. (2013). Indeterminacy and the Small-Improvement Argument. *Utilitas*, 25(4): 433–45.
- Heuer, U. (2011). Beyond Wrong Reasons: The Buck-Passing Account of Value. In M. Brady (ed.), *New Waves in Metaethics*. London: Palgrave Macmillan, pp. 166–84.
- Horty, F. J. (2012). *Reasons as Defaults*. Oxford: Oxford University Press.
- Hieronymi, P. (2005). The Wrong Kind of Reason. *Journal of Philosophy*, 102(9): 437–57.
- Hsieh, N.-h. (2005). Equality, Clumpiness and Incomparability. *Utilitas*, 17(2): 180–204.
- Parfit, D. (1984). *Reasons and Persons*. Oxford: Clarendon Press.
- Parfit, D. (2011). *On What Matters*, vol. 1. Oxford: Oxford University Press.
- Peterson, M. (2007). Parity, Clumpiness and Rational Choice. *Utilitas*, 19(4): 505–13.
- Pritchard, D. (2009). Safety-Based Epistemology: Whither Now? *Journal of Philosophical Research*, 34: 33–45.
- Rabinowicz, W. (2008). Value Relations. *Theoria*, 74(1): 18–49.
- Rabinowicz, W. (2020). Are Probabilities Values? In I. Niiniluoto and S. Pihlström (eds), *Normativity: The 2019 Entretiens of Institut International de Philosophie (Acta Philosophical Fennica, vol. 96)*, Helsinki: Philosophical Society of Finland, pp. 67–99.
- Rabinowicz, W. and Rønnow-Rasmussen, T. (2004). The Strike of the Demon: On Fitting Pro-attitudes and Value. *Ethics*, 113(3): 391–423.
- Rabinowitz, D. (2011). The Safety Condition for Knowledge. *Internet Encyclopedia of Philosophy*. <https://iep.utm.edu/safety-c/>.
- Rønnow-Rasmussen, T. (2011). *Personal Value*. Oxford: Oxford University Press.
- Rowland, R. (2019). *The Normative and the Evaluative: The Buck-Passing Account of Value*. Oxford: Oxford University Press.
- Scanlon, T. M. (1998). *What We Owe to Each Other*. Cambridge: Belknap.
- Schroeder, M. (2010). Value and the Right Kind of Reason. *Oxford Studies in Metaethics*, 5: 25–55.
- Skorupski, J. (2010). *The Domain of Reasons*. Oxford: Oxford University Press.

¹⁹We would like to thank Marta Johansson Werkmäster and Wlodek Rabinowicz for helpful discussions. Thanks also to the participants of the higher seminar at the Department of Philosophy, Lund University, and to the editor and two anonymous reviewers for useful questions and comments. The Swedish Research Council, grant numbers 2018-06698 and 2020-06383, funded this research.

- Sosa, E.** (1999). How to Defeat Opposition to Moore. *Philosophical Perspectives* **13**: 141–54.
- Tappolet, C.** (2013). Evaluative vs. Deontic Concepts. In H. LaFollette (ed.), *International Encyclopedia of Ethics*. Oxford: Wiley Blackwell, pp. 1791–99.
- Whiting, D** (2020). Knowledge, Justification, and (a Sort of) Safe Belief. *Synthese*, **197**(8): 3593–609.
- Williamson, T.** (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.