Care as a Thick Ethical Concept

Ira Chadha-Sridhar[1,2]

Abstract

Philosophers who study care—most often, care ethicists—are involved in an ongoing discussion about the concept of care. Despite the significant progress made in this discussion, certain conflicting images of care seem to persist in the literature. On one hand, as feminist theorists across disciplines have highlighted, care is a complex social practice that is mired in inequality and injustice. The deeply gendered nature of caring and the unequal division of care-work creates and cements structural inequalities. On the other hand, care is also thought of as a moral value or an ideal. The ethics of care—a moral theory with decidedly feminist roots—is predicated on the idea that caring is somehow morally valuable. A discrepancy thus arises: care is a social practice that compounds injustices. But it is also a moral value. What is it about care that makes it malleable to such variations? To pick out this complexity and capture the conceptual nuances at play, this paper suggests that we frame the concept of care as a thick ethical concept. I will first demonstrate why this framing is helpful. Then, I will provide accounts of the descriptive and evaluative elements of the concept of care. I hope to show that it is through a thick conceptual framing that we can make sense of care—both as a complex social practice and an important moral value in the worlds that we inhabit.

Keywords Care ethics · Concept of care · Description · Evaluation · Thick concepts · Feminist philosophy

Introduction

Care plays a crucial role in human life. Most of us will require care at some point in our lives, and often care for others when they need it. To care—and to be cared for—thus seem to be experiences almost universally shared. Despite its near

✉ Ira Chadha-Sridhar
  ic357@cam.ac.uk

1 Faculty of Law, University of Cambridge, The David Williams Building, 10 West Rd, Cambridge CB3 9DZ, UK

2 O.P. Jindal Global University, Sonipat, India

universal presence in our lives though, care varies dramatically with context. It is messy, raw and often deeply personal—taking different forms on the basis of the persons and social groups through which it operates. What, then, does it mean to care? What does the concept of care entail?

Philosophers who study care have undertaken this inquiry and provided us with a range of helpful conceptual accounts (Steyl 2020; Pettersen 2008a, b, 2012; Collins 2015; Engster 2007, 2005; Held 2006; Bubeck 1995; Bowden 1993). Despite these inroads, though, certain tensions about the concept of care seem to persist. In particular, there tend to be different (and potentially conflicting) imageries of care in academic literature. On one hand, care is seen as a social practice that *furthers structural inequality,* particularly along the axis of gender. After all, as feminists have noted, women across the world undertake the majority of care-work and this has cemented structures of male supremacy (Herring 2013, p. 34). On the other hand, though, care is often imagined as an ethical ideal. A prominent school in feminist ethics—the ethics of care—is premised on the idea that caring is somehow *morally valuable*. A conceptual discrepancy thus arises: care is seen as a social site that furthers structural inequality, on one hand, and as a moral value on the other. What is it about care that makes it malleable to such variations?

To answer this, I think we need to ask a related (perhaps prior) question that has so far been overlooked: what kind of concept is the concept of care? I will introduce this question here and wager an answer. In the course of this paper, I will argue that the concept of care should be framed as a *thick ethical concept*. 'Caring' should be thought of as a thick term and the concept of care (to which I refer to in the italicised as *care)* should be framed as a thick ethical concept. Such a conceptual framing can make sense of care—both as a complex social practice and as an important moral value—in the worlds that we inhabit.

I will proceed as follows. In Sect. 'Thick Ethical Concepts' I will briefly introduce thick ethical concepts and discuss their significance in moral philosophy. Very roughly, thick concepts are concepts that combine evaluation and description in ways that their counterpart thin concepts do not (at least not to the same degree or in the same way). In Sect. 'Why Care is a Thick Ethical Concept: Payoffs for Care Ethics' I will demonstrate why *care* should be conceived of as a thick concept. I do so by discussing the payoffs or beneficial upshots of this approach for care ethicists. Finally, in Sect. 'Understanding *Care:* The Descriptive and the Evaluative', I will outline the descriptive and evaluative elements of *care* are and discuss how these elements are related to each other. While our question on conceptual type is indeed analytically distinct from the question of conceptual content, these inquiries are eventually two sides of the same coin.

This paper has implications not just for the philosophical work on *care*, but also for the philosophy of thick concepts. The example of *care* can help clarify how thick evaluation works, demonstrate how different thick concepts work potentially in different ways from one another and make a case for the importance of thick concepts in new areas of moral philosophy. Taken together, this paper seeks to contribute to philosophical inquiry on care ethics and thick concepts—forging a link between these two areas that I hope will be further examined.

Thick Ethical Concepts

In our evaluations of persons and actions, we do not merely evaluate them as right or wrong, correct or incorrect, or good or bad. Rather we may also refer to them as honest, kind, courageous, rude, selfish, selfless, generous and so on. This latter set of concepts describes more about actions than their mere rightness or wrongness. For this reason, such concepts have been called '*thick ethical concepts*'.

The idea behind thick concepts can be traced back to R. M. Hare's 1952 text, 'The Language of Morals'. Here, Hare distinguished primarily evaluative words and secondarily evaluative words which later came to be associated with the terms 'thin' and 'thick' respectively (Hare 1952, pp. 121–122). Other than this, the literature on thick concepts is said to have had two 'founding mothers': Iris Murdoch and Philippa Foot—who articulated the idea behind thick concepts (though they did not use the term) in a seminar in the 1950s. Influenced by their discussion (Abend 2019, p. 210), Bernard Williams introduced thick ethical concepts to mainstream analytic philosophy in 'Ethics and the Limits of Philosophy'.

Williams (1985) introduced certain types of terms or concepts as thick. Unlike thin terms or concepts that merely evaluate something as being good or bad, right or wrong in a thin sense, thick terms and concepts, Williams argued, *combine description and evaluation.* To understand the distinction better, consider this excerpt from Simon Kirchin (2013, pp. 1–2):

> Imagine we are discussing the merits of a mutual friend, Peter. I describe Peter as being good…Describing Peter in this way does not tell us much about his character. There are plenty of ways in which people can be morally good, and plenty of types of action that they typically perform. In contrast, imagine that our discussion continues and I describe Peter as being honest. There are plenty of ways in which people can be honest or show their honesty. However, this more specific description gives us more of a sense of what Peter is like, and perhaps why we think of him as good.

The term 'honest' (and the concept of honesty) refers to a more *specific* trait than the term 'good' (or the concept of goodness). In Kirchin's example, we thus know more about Peter if we know that he is honest than if we know merely that he is good because the term 'honest' indicates specifically that perhaps, in most cases, Peter tells the truth. This is not rendered salient with the usage of the term 'good'. 'Honest' is thicker in its content than 'good', and *honesty* would thus qualify as a thick ethical concept.

While we may be able to intuitively gauge the difference between thick concepts such as *honest* or *courageous*, and thin concepts such as *good* or *bad*, we must probe further to ask: what precisely is the distinction in play? Put simply, a thick ethical concept is a concept that somehow combines evaluation and non-evaluative description (Dancy 2013, p. 44). Through the use of this concept, we evaluate the phenomena (in either a positive or negative manner) and

also describe some features of it in a non-evaluative sense. Thick concepts are thus to be distinguished from purely descriptive concepts and purely evaluative concepts, because thick concepts combine elements of both. Purely descriptive concepts or terms such as '*chair', 'wood', 'dress*' or '*beer*' are used to pick out the features or properties of certain things—to describe them. On the other hand, purely evaluative concepts or terms (thin concepts) such as '*good', 'right', 'bad'* or '*wrong'* are used to evaluate a thing and to indicate some kind of approval or disapproval for the thing in question. By contrast, thick concepts, in Williams's words, do both—they are 'action-guiding' and 'guided by the world' (Williams 1985, p. 144).

Often, the distinction between thin and thick terms and concepts is seen as a distinction of degree rather than a distinction of kind (Scheffler 1987; Hare 1963). On this view, it is not the case that thin concepts contain *no* descriptive content, but rather that they possess *less* of such non-evaluative content as compared to their thick counterparts.[1] Even if we see the distinction between thin and thick as a difference of degree rather than kind, we can agree that there is some intuitive difference at play here that allows us to differentiate broadly between thick and thin concepts. Thick concepts either have a descriptive or non-evaluative element that thin concepts lack, or at least that thin concepts and terms possess to some lesser degree. It is this combination of evaluation and description—or a 'union of fact and value' (Williams 1985, p. 144)—that lends a concept its thickness.

Since the second half of the twentieth century, thick terms and concepts have become an important theme in different areas in moral, political and legal philosophy (Kirchin 2013, 2017; Zangwill 1995; Putnam 2002; Enoch and Toh 2013). By drawing care ethicists' attention to this area of scholarship at the intersection of philosophy of language and analytic moral philosophy, I aim to show that there are very good reasons to hold that *care*, in fact, should be framed as a thick ethical concept.

Why *Care* is a Thick Ethical Concept: Payoffs for Care Ethics

Let us begin by considering how we use the term 'caring' in ordinary language. Suppose X and Y are two friends who live together. X falls suddenly ill and cannot make a meal for themselves, and so, Y steps in by preparing a meal for X. We may thus refer to Y's act of cooking a meal for their friend as 'caring'. Now suppose this is not a one-time occurrence but a matter of routine. X and Y both often cook for one another when either one is unable (or even unwilling) to. We may call their relationship a 'caring' one and refer to each of them, in turn, as 'caring' people. Such statements that use the term 'caring' *describe* certain facts about X and Y. But there is also an *evaluation* that accompanies this description. We evaluate X

---

[1] For instance, Hare argues that the non-evaluative or descriptive meanings of thick terms are 'more firmly attached' to them than their evaluative counterparts, while the converse is true for thin terms (Hare, *Freedom and reason,* pp. 24–25).

and Y's actions towards one another here as being good or morally praiseworthy acts. Statements that use the term 'caring' thus walk such a fine line between the descriptive and evaluative.

Thick concepts are characterised specifically by the fact that they combine description and evaluation. Even at first glance, then, *care* seems to bear a resemblance with such thick ethical concepts. Consider two cases of paradigmatical thick concepts: *selfishness* and *courage*. When we call someone selfish, we *describe* the fact that they perhaps unthinkingly put themselves over others, apart from *evaluating* this action as being bad or wrong. Or when we call someone courageous, we seemingly *describe* the fact that they are not deterred by danger or pain, and also *evaluate* this quality as being admirable, worthy of respect or generally good. Similarly, when we refer to a person or an action as 'caring', we *describe* something about the way that they are and *also evaluate* this as being generally good, desirable or right.

Consider the following excerpt from Held (2006, p. 38):

> If we say of someone that ''he is a caring person,'' this includes an evaluation that he has a characteristic that, other things being equal, is morally admirable…It is highly useful to be able to characterize people (and societies) in specific and subtle ways, recognizing the elements of our claims that are *empirically descriptive* and those that are *normative*…Caring thus picks out a more specific value to be found in persons' and societies' characteristics than merely finding them to be good or bad, or morally admirable or not.

From this excerpt we see that the phenomenon of care functions in a very similar way to the phenomena that thick ethical concepts seek to pick out. On the one hand, descriptions of care are accompanied by (generally positive) evaluations. On the other, care is not like thin ethical concepts such as good and right: it is more *specific*. Care thus has both this descriptive element and a connection with moral value. In order to capture this subtlety, I suggest that we frame *care* as a thick concept to pick out the evaluative dimension of care without compromising on its descriptive distinctiveness.

I will now elaborate upon certain *payoffs* of framing *care* as a thick ethical concept in the following parts Care as Practice, Care as Value, Care Ethics as a Moral Theory: From Is to Ought, Particularity and Care and 'Good' Care.[2] We will see that each of these payoffs give us reason to think of *care* in a thick conceptual sense and motivates a study of care in these terms.

Care as Practice, Care as Value

First, framing *care* as a thick concept helps reconcile two main features of caring. While care is a set of actions, or a social practice,[3] it is also a moral value. This

---

[2] I borrow the helpful language of 'payoffs' from David Enoch and Kevin Toh (2013).

[3] For distinctions between caring actions, activities and practices, see Steyl (2020). On practices generally, see Sangiovanni (2016).

duality, between the social and the moral, was first stressed by Virginia Held. While Held argued that to care for another is a social practice—a series of doings or actions that are geared towards meeting the needs of others—she resisted the characterisation of care as being just this and cautioned against constructing *care* in entirely non-normative terms (Held 2006, pp. 36–37).

Importantly, care ethicists following from Held have also stressed the importance of viewing *care* in these decidedly normative terms. This is crucial considering the historical devaluation of care-work along structural lines (Chatzidakis et al. 2020, p. 24). The majority of care-work has been performed by women, or people of colour, or people from marginalised caste and ethnic groups in a given society (or perhaps the marginalisation of these groups has been legitimised *because* they participated in care-work). In light of this devaluation, care ethicists have struggled to get the social practice of care recognised normatively. With this in mind, care ethicists emphasise the duality of value and practice in *care*. Consider Joan Tronto and Bernice Fisher who define care, itself, as being geared towards morally valuable ends. They write:

> Caring is…"everything that we do to maintain, continue and repair our world so that we may live in it as well as possible." (Tronto 1993, p. 103; Fisher and Tronto 1990, p. 40)

In a similar vein, Sara Ruddick points out that three meanings of the term 'care' have emerged: it is either a form of labour, or a relationship or a kind of ethic (Ruddick 1998, p. 4). While the first and second meanings deal with the social phenomenon of care, the third meaning hints at moral value. For Engster, the social and the moral aspects of *care* are sometimes even at odds with one another in that they have different emphases. He writes (Engster 2007, p. 21):

> Caring for human beings, in turn, is usually defined as *either* virtue or practice. When caring is defined as a virtue, the focus is on inner traits, dispositions, and motivations of the caring person …When caring is defined as practice, by contrast, the focus is on a person's external actions and their consequences… These two definitions of caring are not entirely opposed… The two approaches nonetheless have different *emphases*.

While the idea that *care* picks out both a social practice (understood here as a series of actions geared towards a common purpose) and a moral value has thus been reflected—in different forms—across the literature in care ethics, very little has been said about *how* and in what ways practice and value are interrelated. Held asserts that the practice of care *incorporates moral value* and that caregiving has values *embedded* within it (Held 2006, p. 38). But it is not clear what such a claim about embedment entails. Is moral value embedded in every act of caring? Or are we thinking of caring actions and the value of care as different things that may only sometimes overlap?

I think that the relationship between practice and value can be best picked out through a thick conceptual frame. As we have seen, thick concepts combine description and evaluation. We can both pick out *care*'s descriptive element—by

describing caring actions or practices—and identify how our descriptions of such actions are often accompanied by an *evaluation*. Further, there is extensive material in the literature on thick ethical concepts on how the descriptive and evaluative elements are combined within a concept. This material indicates that it is not the case that evaluations accompanying descriptions of *care* must be necessarily positive. Rather, this evaluation may vary, in particular ways, with context (Vayrynen 2011). Care ethicists have not so far paid attention to this material or made use of the findings in this field. But this material can greatly inform how we view the relationship between the activity of caring and moral value. I will discuss both these elements further in the subsequent section of this paper. For now, however, we can note that the material in the work on thick concepts can lend a frame with which to understand how practice and value are interrelated in one concept.

It may be argued, at this juncture, that we could employ different concepts of care—one which picks out care as a practice and the other that picks out care as a value—instead of striving for a unitary concept. But that would distort the way we use the term 'care' usually in everyday discourse. If I say that a certain phenomenon X is 'caring', this is not similar to a description of a certain phenomenon as being 'red', for instance. The term 'care' or 'caring' is typically loaded heavily with an evaluative dimension—wherein the usage of such terms indicates that the speaker is making some kind of evaluation that accompanies their description. This close relationship between practice and value makes it difficult (and indeed undesirable) to study these two aspects as separate phenomena. If we do so, we miss the close interrelationship between the two: after all the value arises from the practice, and the practice cannot be understood without an understanding of how it is valued in the world around us.

In light of this, rather than resorting to a concept of care that focusses on a single aspect of what it means to care, we require a conceptual apparatus that does both: picks out social fact and moral value. As thick concepts are concepts that do precisely this, they make for a compelling conceptual frame for *care*.

Care Ethics as a Moral Theory: From Is To Ought

Second, and in relation to our discussion above, the *ethics of care*—the moral theory that is based on care—has grappled with a tension between the descriptive and the normative. Care ethicists have all focussed (in some way or the other) on care's moral value and potential—arguing that care ethics can offer 'a new orientation to social and political thinking' (Engster and Hammington 2015, p. 1). Care ethics is not interested, as Held notes, merely in describing the practices of care as they have evolved across history and cultures (Held 2006, pp. 37–38). Rather it is an ethical theory that is based on the idea that care is importantly associated with moral value and it requires evaluations and judgements, not just empirical findings (Held 2014, p. 114).

But as Pettersen writes, if this is the case, care philosophers who are committed to building care ethics must pay closer attention to 'the interface between the normative

and the empirical' (Pettersen 2008a, b, p. 61). Care ethics has often been criticised for being unsuccessful in bridging this very gap. I think that the fact that this has been a challenge for care ethics is also perhaps to do with historical reasons. Carol Gilligan's early work (Gilligan 1982)—which provided the foundation of the care ethical tradition—was in the field of psychology. Petterson (2008a, p. 37) writes that this has sometimes posed a problem:

> It appears as if Gilligan draws normative conclusions from her psychological thesis on human nature, thus blurring the distinction between 'is' and 'ought'.

Because Gilligan's work was the bedrock for care theory, care ethicists have had to negotiate how to think of *care* as a moral concept that can provide a foundation for a moral theory. This has led to tensions in care ethics—fractures between the descriptive and the evaluative, between the empirical and the normative, between fact and value.

I think that framing *care* as a thick concept helps think through these tensions and provides ways in which care ethicists may satisfy their justificatory burdens. Drawing from literature on thick concepts—that I will discuss in significant detail in Sect. 'Separability: A Final Payoff'—there seem to be two ways to think about the relationship between *care*'s descriptive and evaluative components.

First, we may view the descriptive and evaluative elements as inseparable in the sense that the social fact of *care entails moral value*. On this view, *care* in and of itself—like other thick concepts—bridges the is–ought gap and provides possible ways out of the naturalistic fallacy. The is–ought gap, attributed to David Hume (1739), holds that we cannot validly arrive at evaluative statements from a set of purely descriptive statements. If the descriptive and evaluative elements of a concept are inseparable, thick concepts can potentially bridge this gap. As Philippa Foot observes, the description of certain terms such as 'rude' which could be described as 'causing offence by indicating a lack of respect' *contains* negative evaluation (Foot 1958). If we think of *care* in these terms, where the descriptive and evaluative components are inseparable, we arrive at one way of bridging the is–ought gap.

Second, in the alternative, we may view the relationship between the descriptive and evaluative component as being fundamentally separable. On this view, describing an action as 'caring' does not entail that the action has moral value. If this view is adopted, care theorists will have to look outside *care* to bridge the is–ought gap. They will have to rely on moral principles outside the conceptual framework of care: such as the principle of fairness or reciprocity.

Depending on how we view the relationship between *care*'s components, we can gauge whether care itself can act as a bridge between 'is' and 'ought', or whether we require moral principles external to *care* to bridge this crucial gap. This discussion could be particularly important in the context of Steven Steyl's recent and valuable efforts to develop a care ethical theory of right action (Steyl 2021). Steyl suggests that care theorists must work towards building such a theory through available conceptual resources. Framing *care* as a thick concept and determining whether a separable, or inseparable conceptual account of *care* should be adopted, could dramatically strengthen the conceptual resources available to care ethicists.

Particularity

Third, another reason why care ethicists should frame *care* as a thick concept is that thick concepts have been considered to be more particular and grounded in social context than their thin counterparts. There is a layer of descriptive specificity that comes with thick concepts, in a way that thin concepts lack (at least to the same degree). To elaborate on this point, consider this extract from Scanlon (2003, pp. 276–277):

> [T]he thinness of thin concepts such as right, wrong, duty and obligation lies first and foremost in the abstractness, hence relative emptiness, of the ethical ideas that they involve, by contrast with concepts such as coward, lie, betrayal, brutality, honor, and gratitude, whose greater content reflects the *distinctive character of particular social worlds*.

Care ethicists have often resisted characterising *care* through thin value concepts. Held (2006, p. 38) writes that we cannot think of *care* in the 'most general value terms' such as good or right. Rather, to care for someone has a layer of specifity that more thin concepts lack. In light of this, one of the most important features of care ethics has been its particularity. Steven Steyl (2020, p. 279) summarises this commitment to particularity aptly, when he says:

> Care ethics has always regarded itself as an ethic at the coal face. Care ethics, it is said, is a *practical* ethic, rooted in the lived experiences of caregivers and care-recipients. For this reason, care ethicists have often taken some type of action, activity, work, labour, or practice to be a, if not *the*, central moral concept.

Care ethics is thus rooted resolutely in 'particular social worlds' and takes certain activities or practices to be central moral concepts. But how can social practices, often mired with inequalities and injustices, be useful moral concepts? The varied practices of care, the changing meanings of care over time (Tronto 2013, pp. 1–2, 18–20), the unequal distribution of care-work across societies and the fact that most care-work continues to be undertaken by women (Alonso et al 2019; Herring 2013, p. 34), particularly those from marginalised communities (Raghuram 2021, p. 617), shows that care ethicists must think further about how and why the social phenomenon of care is at all morally valuable. Framing *care* as a thick concept provides a solution to this problem. The literature on thick concepts—particularly the wealth of recent work on the relationship between evaluation and description—can help clarify how certain social facts are related with moral value.

Apart from the internal value of this move for care ethics, a thick conceptual framing also helps *locate* care ethics within the landscape of analytic moral philosophy. Stephanie Collins writes that despite the amount of recent work on it, care ethics continues to exist as 'a somewhat shadowy entity at the outskirts of debates in analytic moral philosophy' (Collins 2015, p. 2). If *care* is constructed as a thick concept, care theory acquires a closer affinity with contextual particulars (as compared to a theory that has thin value concepts at its core). This affinity with

particularity can pave the way for framing care ethics, potentially and after more research on this subject, as a type of moral particularism (Dancy 1992). Seeing care ethics as a *type of particularism* would resolutely locate care ethics in the domain of analytic moral philosophy and facilitate further engagement with the approach and its core tenets.

Care and 'Good' Care

Fourth, in recent discussions on *care,* some philosophers have proceeded with the understanding that we should think of care in two ways: as care simpliciter and as 'good' care. Collins (2015) distinguishes 'care' (caring actions and attitudes) from what she calls 'morally valuable care'. Steyl (2020, pp. 289–290) also presents an account of caring actions wherein he argues that 'an action that is caring is not, *eo ipso*, morally good or right'. Pettersen seems to assume this distinction as well when she argues that Gilligan's theory must be defended further—'since what is to be proven is precisely *which* descriptions (of care) may be recognised as having normative value' (Pettersen 2008a, b, p. 40). There thus seems to be an emerging distinction between care simpliciter and 'good' care, which has become a prominent feature in contemporary scholarship.

If we are to operate with this distinction (which I think we should), we require the conceptual tools to think through *how* care is related with good care. Are we to think of care and good care as distinct concepts? Or are we to admit that they are tied together in some way—that may not make the link between them necessary, but still shows that they are often rather closely tied together? I think the philosophical investigations on thick concepts can, once again, help us immensely in thinking through this distinction.

As I have stated earlier, I do not think that having *care* and 'good' care work as *two distinct concepts* is a viable option. We do not use the terms 'care' and 'good care' as entirely separate ideas in our ordinary language, cultures and conversations. Rather, the evaluation of care is intimately tied with its description. Considering thick concepts combine fact and value within a single conceptual structure, they are best suited to pick out this complexity. Furthermore, as I will demonstrate in Sect. 'Evaluation', most philosophers who work on thick concepts believe that while thick concepts have a *default valency,* this valence may vary depending on context (Vayrynen 2011; Kirchin 2007a, b; Norman 2007). A thick conceptual apparatus can help us make sense of this variation and shed light on emerging distinctions in care ethical scholarship between care and 'good' care.

Cumulatively then, I have shown that constructing *care* as a thick concept not only serves a clarificatory purpose, but also opens up new and important areas of inquiry.

Understanding *Care:* The Descriptive and the Evaluative

Having shown why this conceptual frame is helpful, we are now well placed to think closely about what *care* would entail. I will outline both elements of *care*: the descriptive and the evaluative element in Part Description and Evaluation respectively. Note that I will proceed with this discussion with the assumption that the descriptive and the evaluative elements of care can be thought of as distinct, at least for the purposes of analytic clarity. I adopt this stance merely to elucidate that *care* has these two elements and am not committed to the position that these elements are at all separable from each other, or indeed that they may be neatly 'disentangled' (Blackburn 2013). As I will show in Part Separability: A Final Payoff, the question of separability—which is arguably the most important meta-ethical debate on thick concepts—is a difficult and important issue, that I do not seek to take a definite position on in this paper. My goal is merely to show that *care* is in fact a thick concept and that the literature on separability can help nuance our understandings of *care* in ways that are beneficial for care theory.

Description

Let us take a look at the descriptive element of *care*. The first thing to ask here is whether and how *care* is distinct from more general (indeed thin) evaluative terms such as 'good' or 'right'. It seems like it is. Several people, relationships and interactions may be considered to be good, but only some of these can correctly be called 'caring'. Consider the difference between a good relationship and a caring one. Not all good relationships are caring: the former does not entail the latter. I may have a good relationship with my colleague in a workplace, where we are both cordial with one another and stay out of each other's way. We may not know any particular information about each other's specific needs, personalities, preferences or interests, and may do nothing to make each other feel better on a daily basis. But we work well together professionally in all our interactions and greet each other politely in formal settings. While our relationship may be called a *good relationship,* it is certainly not a caring one. The term 'care' (and thus *care* as a concept) has some descriptive element that adds a layer of specificity to it—one that may be entailed by the term 'good' in certain contexts, but is certainly not spelled out merely through the usage of the term 'good'.

I will focus here on caring actions—what it means to 'care for', 'take care of' or 'give care to' another person through the performance of some actions.[4] There

---

[4] The term 'care' can refer to a few distinct (though related) phenomena. We might 'care about' some-one or something. This 'caring about' aptly describes care as an attitude. In contrast, when we 'care for' or 'take care of' someone (or perhaps something), we perform certain *actions* (that may or may not be coupled with certain desirable attitudes) towards them.

Care - Pro	Care – Neutral	Care - Con

Fig. 1 Evaluative stances

is extensive literature on what caring actions entail (See Steyl 2020) and it is not my aim to delve too deeply into the literature. However, even a brief glance at the literature, shows that caring actions are those that *meet the needs of another being* (Randall 2020; Held 2006; Steyl 2020; Miller 2012; Kittay 1999). As Tronto (1993, p. 105) writes, what is definitive about care 'seems to be a perspective of taking the other's needs as the starting point for what must be done'. Even Steyl (2020, p. 285) shows that need is a universal element across most competing definitions of care and concludes that one does not indeed have to endorse 'a full definition of care in order to admit that some elements of care'—including the idea of meeting needs— are necessary elements in our conceptual account.[5]

Therefore, to perform a caring action is to act in order to meet someone's needs. When we provide food to a sick person, bathe an infant or tend to someone with a fractured leg or arm, we are recognising some need that they have and acting in order to meet this need. This is what, we can agree, makes an action caring. I therefore suggest a very rough formulation to outline the descriptive element of *care:*

> D: A has cared for B by Φ-ing, if by Φ-ing, A intends to meet any of B's needs or go some way towards meeting B's needs.

We can say that this specific act—of acting with the intent to meet the needs of another—is what distinguishes 'caring actions' from merely 'good' ones. It is this aspect that lends *care* its descriptive distinctiveness.

Evaluation

While the practice of care can be explained through a rough formulation such as D as we have identified above, D does not capture the evaluation that accompanies descriptions of care. A deeper glance into the subject of evaluation leads us, in particular, to one very important question. Is caring always positively evaluated? Or can we think of negative evaluations associated with caring? The literature on thick concepts can help by introducing—and potentially answering—this question.

Most philosophers who study thick concepts seem to agree that some degree of *contextual variability with respect to evaluative valence* exists in the case of most (if not all) thick ethical concepts. Even if courage and honesty are typically evaluated favourably, there are certain cases where a negative evaluation would still be expected and warranted (Vayrynen 2011). I think the same is true of care as well. While it is the case that caring actions are *mostly* evaluated positively, or that there

---

[5] Collins (2015) uses the term 'interests' instead of 'needs'.

is a pro-attitude to these evaluations, there are also certainly exceptions to this case that can be brought to light.

If this is indeed true, it becomes important to ask *how* this evaluation varies on the basis of context. I will now argue that a close look at *care* shows that we can have (at least) three types of evaluations of caring actions. These stances can be summarised as follows:

The *Care*—Pro stance in Fig. 1 above is perhaps the default valence associated with care.[6] Most often, when we describe some action as being 'caring', there is a positive valence associated with it. To care is considered to be a good thing, and when we say of someone that they are 'caring', we mean this oftentimes as a positive, valuable trait—even a virtue (Slote 2001, 2007; Halwani 2003; Steyl 2019). A relationship in which both people help meet the needs of another is considered to be an example of a good relationship, as much as it is an example of a caring one. Predominantly, then, I think that a positive evaluation accompanies descriptions of caring actions.

On the other hand, what I have labelled as the *Care*—*Con* stance is perhaps the least intuitive stance of the three. But it is still not inconceivable. There may be actions that are 'caring' but would still be considered to be morally 'bad' or wrong, *all things considered*. For instance, when severe large-scale war crimes have been committed, there have often been a group of doctors working along with those directly conducting mass murder or genocide (Mehring 2011). Suppose those doctors were charged with the job of attending to the medical needs of the perpetrators of genocide, in such a manner and with such an intent as to allow the perpetrators to go about their actions of mass murder or torture efficiently. In this case, then, the doctors—attending to basic medical needs–may be said to be involved in 'caring' actions. But these caring actions would almost certainly be evaluated negatively, for their role in furthering an unjust scheme of crimes against humanity. This shows us that in certain circumstances, caring actions may be evaluated *negatively* by a community of observers.[7] It is not necessarily the case that 'to care' is to perform a morally praiseworthy action.

At the middle of these two ends, we have what I have termed as the *Care*—Neutral stance. Under such an approach, our evaluations of caring actions may not decisively conclude with either a *pro* or *con* valence and may be more complex. There may be cases wherein our evaluations of caring actions may be inconclusive—they are neither pro nor con, all things considered. For instance, if A and B are in a relationship wherein the burden of care-work falls solely on A, we are faced with a complex case of care. Suppose further that A gets a sense of themselves out

---

[6] On whether thick concepts necessarily possess a default valence, see Kirchin 2013, 2007a, b; Norman 2007.

[7] Alternatively, we may hold that these kinds of caring actions are still good in themselves, but they merely contribute to a larger scheme that is unjust. On such a view, care is evaluated on the basis only of how it achieves certain goals, in certain ways. The larger social scheme of which it is a part is immaterial. But as critics of care ethics often note, it is important that our understandings of *care* allow us to resist the complicity of care in projects of large-scale evil. There is a difference I want to hold on to between care that is, and care that is not, 'complicit in evil-doing', as Claudia Card notes. It is for this reason that I have included the 'scheme' of the action in the subsequent part of the paper and want to point to all-things-considered evaluations of care on the basis of the social scheme it contributes to.

of doing all the care-work and enjoys doing this work most of the time. We may have reason to evaluate A's caring actions as being good because they help cater to B's needs. But equally, we would also be worried about the unequal nature of the relationship between A and B. A non-reciprocal relationship (particularly if it is further complicated on a power axis of the likes of gender), would be a complex evaluative case.

These facts would complicate our evaluations of A's actions. We may thus conclude that we have reasons to evaluate A's actions both negatively and positively and may thus not settle neatly on a pro–con valence (Dancy 2013). Note that this case, and all cases that I can think of as falling under the *Care*—NEUTRAL stance, are not purely descriptive cases. We *do not merely describe* the actions here as caring. There is still very much an evaluative stance that accompanies the description and thus *care* here is still a thick concept. The only point is that this evaluative stance cannot neatly be categorised as either pro or con.

Having established that evaluation of caring actions can broadly be of three types, we can see that there is a spectrum of evaluations upon which caring actions rest. We often hear statements such as 'P took care of Q. P is such a good friend'! Or 'P cares for Q despite Q not caring for P. Their relationship is unequal'. This shows that *care*'s valence varies with certain factors. But what are these factors? On what basis do we make such evaluations of caring actions?

I think our evaluations of caring actions depend on how caring actions fare on the following three metrics: first, the *consequence* of the action; second, the *process* by which needs were met; third, whether the action contributed to a larger *scheme* of injustice. These three factors—the *consequence, process* and *scheme*—are at the heart of our evaluation. Depending on each of these factors (and often a complex combination of these three), we assess a caring action as being morally praiseworthy or not.

Let us take a closer look. On formulation D stated earlier, a caring action is one that is done with *intent* to meet someone's needs. Whether or not the needs are actually met, or whether the care is 'completed' as Eva Kittay (2014, pp. 33–34) calls it, is not important at the stage of description. I think this factor of the *consequence* of the caring action—or how successful it was in meeting the needs of the recipient—matters now in our *evaluations* of caring actions. Think of a case wherein A, B and C all set out to meet D's needs. A meets these needs. B is unsuccessful in meeting them. C's actions are not only unsuccessful in meeting D's needs, but are also actually harmful to D. All other things being constant, I think our *evaluations of the care* given in each case would *vary* on the basis of the outcome— we may be comfortable saying that D is better cared for by A than by B, and better cared for by B in comparison to C.

Similarly, if we think now about the *process* by which needs are met, we see that this too influences our evaluations of caring actions. Suppose I am very ill and request two of my friends—P and Q—to bring me a meal. Both bring me meals. However, P brings this meal grudgingly and is rude and resentful towards me. Q gives me the meal in a respectful manner and is kind, attentive and responsive to my needs. I (and indeed I suspect most of us) would be tempted to evaluate Q's actions

more favourably than we evaluative P's. This is because the *process* by which needs are met is important.

Engster (2007) suggests that the values associated with the process of caring are attentiveness, responsiveness and respect—which all feature also in Tronto's work on the moral qualities (or 'phases of care') that are associated with caring (Tronto 1993, 2013).[8] I think this is a fair characterisation to work with here (though I do not seek to justify here that these are the only, or main, values associated with caring).[9] Drawing from Engster's account, we may say that we are more likely to evaluate a caring action favourably if the agent has been attentive, responsive and respectful to the recipient's needs when performing said action.

Finally, our evaluations of caring actions, all things considered, also depend on the larger scheme to which such actions contribute. Steyl discusses this in the following excerpt:

> As Tronto points out, care's social nature introduces a set of external power dynamics to caring relations (Tronto 2010; 2013). One thought underpinning her work, and the work of care ethicists like Held and Fiona Robinson, is that caring relations are subject to authorities beyond the caregiver–recipient relation (Robinson 1999; Held 2006, 2011; Robinson 2011). Caring relations are caught up in the power dynamics of the wider family, community, profession, state, and even the international community.[10]

The larger scheme under which the caring action is performed—or the social practice of which it is a part—matters for our evaluations of caring actions. Feminist philosophers have highlighted the extent to which care has been a burden for women, and has been a tool in furthering patriarchal oppression.[11] If caring actions seek to contribute to, or even uphold, a larger scheme of injustice within the family or community, we may again have reason to evaluate said action negatively. For instance, if a woman bears the sole burden of care-work for an abusive husband, and maintains an abusive relationship with him, we would certainly not think of her caring actions as being 'morally praiseworthy' all things considered. Her actions in this case further her own oppression and in some cases may expose her to direct physical harm. Thus, a more complex evaluative stance (something like *Care–Neutral* or a *Care–Con* stance) would be adopted to evaluate the woman's caring actions. Thus, our all-things-considered evaluations of caring actions will certainly depend, in important ways, on the scheme or practice of which caring actions are a part.

Therefore, the consequences, the outcome and the scheme of which the caring action is a part would all contribute towards determining whether we would have

---

[8] Note that Engster errs because he *describes* caring actions in this way and does not distinguish between 'care' and good care'.

[9] Tronto's discussion (2014, pp. 34–35) highlights responsibility and competence along with attentiveness, responsiveness and respect/trust or solidarity as the five ideals that correspond with respective phases of care.

[10] Steyl, pp. 292–293.

[11] Herring, *Law and the relational self,* p. 53.

a *Care*—Pro, *Care*—*Con* or *Care*—Neutral evaluative stance to a certain caring action.

Separability: A Final Payoff

Lastly, the debate between the separationists and the non-separationists is perhaps the most discussed question among philosophers who study thick concepts. As I briefly mentioned in Sect. 'Care Ethics as a Moral Theory: From Is to Ought', the separationists think that the descriptive and evaluative components can be separated from one another, while the non-separationists think that they cannot (Kirchin 2013). I do not seek to defend either the non-separationist or the separationist position here—this is an important discussion for a later juncture. I am merely interested in showing that they are important lines of inquiry.

If care ethicists engage with this literature on separability, we may be able to think more closely about how exactly our descriptions of caring actions are held together with our evaluations of the practice. Can these two components be disentangled from one another, at least in theory? Or is there a necessary overlap that makes such disentangling impossible? Consider Kirchin (2013, pp. 11–12), who writes:

> Imagine an anthropologist going into a society and trying to pick out what a certain evaluative concept stands for. As part of this, or perhaps first of all, she will have to work out what the concepts are that are in use, which ones are evaluative, and so on… Perhaps the outsider encounters a friendly insider who points out which collection of descriptively characterized things are kind and which are not. But, after the insider leaves the outsider, will she be able to continue and apply her knowledge to descriptively new cases? [No] … People can see what various descriptively characterized collections have in common *only if they can latch onto the evaluative point* of putting these items into a collection in the first place.

This is of deep interest in the context of care. Can anthropologists who study care, for instance, make sense of the concept of care without sharing its evaluative point? Or is it possible to identify phenomena as 'caring' only if we do indeed 'latch on' to the concept's evaluative point, as the non-separationists claim?

If a non-separationist view is taken, to call something 'caring' is to generally evaluate it favourably and this may be good enough reason to recommend it or advocate for it. The recommendations made by care ethicists follow directly from the classification of an action as caring. On the other hand, we could also potentially endorse a separationist account of *care*. We could say that D is separate from the evaluative component that accompanies D. It is merely attached to D. On this account, for *care* to be reason-giving and provide sound reasons for action, there must be something else in the picture. What, then, is this something else?

This opens up another debate in care ethics—the debate on moral principles. There is an ongoing discussion among care theorists about whether principles are

necessary for care ethics (Stark 2010; Collins 2015, pp. 17–33) for the approach to navigate the distance between the descriptive and the normative. Depending on the position we take on the separability question, we can establish whether we require external moral principles to help navigate the distance between the descriptive and the normative—or whether *care* itself is a thick concept that can navigate this distance. Closer engagement with the literature on separability and evaluation is thus very much warranted.

Conclusion

In his illuminating work on thick concepts, Kirchin writes, 'thick concepts are important to us and our world because they seem to be a necessary way of understanding what the world and its people are' (Kirchin 2013). I think framing *care* as a thick concept is necessary to make sense of how it seems to work—both as a complex social practice, which often furthers inequality, and as a thing of moral value. It is through such a thick framing that we can make sense of what it means to care, both in our social worlds and moral imaginations.

Cumulatively, I outlined *five payoffs* of framing *care* as a thick concept (though I am sure there are more). *First*, a thick conceptual structure reconciles two aspects of care—care as practice and care as value—and *second*, it helps care ethicists navigate ways out of the is–ought gap to make theirs a more robust moral theory. *Third*, thick concepts are rooted in particular social worlds and connected to natural facts to a greater degree than their thin counterparts. This makes for the ideal conceptual backdrop for care ethicists, who have largely been sceptical of abstraction, and helps locate care ethics—potentially as a species of particularism—within the terrain of analytic moral philosophy. *Fourth,* framing *care* as a thick concept can help distinguish clearly between care simpliciter and morally valuable care. This brings us to the *fifth* and final payoff I have outlined: a thick conceptual framing uniquely illuminates *care*'s distinctive conceptual content. I have shown this by demonstrating that *care*'s descriptive element can be captured through something like formulation D, and then discussed how *care*'s evaluative element works. I demarcated three evaluative stances that we may take when we describe something as caring: CARE—PRO, CARE—CON AND CARE—NEUTRAL stances and outlined the *factors* upon which our evaluation depends: the *consequence* or *outcome,* the *process* and the *scheme* of which the action is a part. I finally briefly tied this up to the discussion on separability and demonstrated the importance of engaging further with it in future scholarship.

In this article I have argued that if we frame *care* as a thick ethical concept, there is much that we stand to gain. A thick conceptual framing may not be able to settle the philosophical debate on what it means to care. But it certainly brings us closer to an answer.

Declarations

References

Abend, Gabriel. 2019. Thick concepts and sociological research. *Sociological Research* 37 (3): 209–223.

Alonso, Cristia et al. 2019. *Reducing and redistributing unpaid work: Stronger policies to support gender equality*. International Monetary Fund Working Paper. WP/19/225.

Blackburn, Simon. 2013. Disentangling disentangling. In *Thick concepts*, ed. Simon Kirchin. Oxford: Oxford University Press.

Bowden, Peta. 1993. *Caring: Gender-sensitive ethics*. London: Routledge.

Bubeck, Diemut. 1995. *Care, gender, and justice Oxford*. Oxford: Oxford University Press.

Chatzidakis, Andreas, et al. 2020. *The care manifesto: The politics of interdependence*. Verso Books.

Collins, Stephanie. 2015. *The core of care ethics*. London: Palgrave.

Dancy, Jonathan. 1992. Caring about justice. *Philosophy* 67 (262): 447–366.

Dancy, Jonathan. 2013. Practical concepts. In *Thick concepts*, ed. Simon Kirchin. Oxford: Oxford University Press.

Engster, Daniel. 2005. Rethinking care theory: The practice of caring and the obligation to care. *Hypatia* 20 (3): 50–74.

Engster, Daniel. 2007. *The heart of justice: Care ethics and political theory*. Oxford: Oxford University Press.

Engster, Daniel, and Maurice Hamington, eds. 2015. *Care ethics and political theory*. Oxford: Oxford University Press.

Enoch, David, and Kevin Toh. 2013. Legal as a thick concept. In *Philosophical foundations of the nature of law*, ed. Wil Waluchow and Stefan Sciaraffa. Oxford: Oxford University Press.

Eva Feder Kittay. 2014. The completion of care—With implications for a duty to receive care graciously. In *Care professions and globalization*, ed. Ana Marta González and Craig Iffland. New York: Palgrave Macmillan.

Fisher, Berenice, and Joan Tronto. 1990. Toward a feminist theory of caring. In *Circles of care: Work and identity in women's lives*, eds. Emily K. Abel and Margaret K. Nelson. Albany, NY: State University of New York Press.

Foot, Philippa. 1958. Moral arguments. *Mind* 67 (268): 502–513.

Gilligan, Carol. 1982. *In a different voice: Psychological theory and women's development*. Cambridge, MA: Harvard University Press.

Halwani, Raja. 2003. Care ethics and virtue ethics. *Hypatia* 18 (3): 161–192.

Hare, R. M. 1952. *The language of morals*, 1st edn. Oxford: Clarendon.

Hare, R.M. 1963. *Freedom and reason*. Oxford: Oxford University Press.

Held, Virginia. 2006. *The ethics of care: Personal, political, global*. New York: Oxford University Press.

Held, Virginia. 2011. Morality, care, and international law. *Ethics and Global Politics* 4 (3): 173–194.

Held, Virginia. 2014. The ethics of care as normative guidance: Comment on Gilligan. *Journal of Social Philosophy* 45 (1): 107–115.

Herring, Jonathan. 2013. *Caring and the law*. Oxford: Hart Publishing.

Hume, David. 1739/ 1958. *A treatise of human nature*, ed. L. A. Selby-Bigge. Oxford: Oxford University Press.

Kirchin, Simon. 2007a. Moral particularism: An introduction. *Journal of Moral Philosophy* 4 (1): 8–15.

Kirchin, Simon. 2007b. Particularism and default valency. *Journal of Moral Philosophy* 4 (1): 16–32.

Kirchin, Simon. 2013. Introduction: Thick and thin concepts. In *Thick concepts*, ed. Simon Kirchin. Oxford: Oxford University Press.

Kirchin, Simon. 2017. *Thick evaluation*. Oxford: Oxford University Press.

Kittay, Eva. 1999. *Love's labor: Essays on women, equality, and dependency*. New York: Routledge.

Miller, Sarah Clark. 2012. *The ethics of need: Agency, dignity and obligation*. New York: Routledge.

Norman, Richard. 2007. Particularism and reasons: A reply to Kirchin. *Journal of Moral Philosophy* 4 (1): 33–39.

Pettersen, Tove. 2008a. *Comprehending care: Problems and possibilities in the ethics of care*. Lanham, MD: Lexington Books.

Pettersen, Tove. 2008b. The ethics of care: Normative structures and empirical implications. *Health Care Analysis* 19 (1): 51–64.

Pettersen, Tove. 2012. Conceptions of care: Altruism, feminism, and mature care. *Hypatia* 27 (2): 366–389.

Putnam, Hilary. 2002. *The collapse of the fact/value dichotomy and other essays*. Cambridge, MA: Harvard University Press.

Raghuram, Parvati. 2021. Race and feminist care ethics: Intersectionality as method. In *The changing ethos of human rights*, ed. Hoda Mahmoudi, Alison Brysk and Kate Seaman. Massachusetts: Edward Elgar Publishing.

Randall, Thomas E. 2020. Justifying partiality in care ethics. *Res Publica* 26 (1): 67–87.

Robinson, Fiona. 1999. *Globalizing care: Ethics, feminist theory, and international relations*. New York: Routledge.

Robinson, Fiona. 2011. *The ethics of care: A feminist approach to human security*. Philadelphia, PA: Temple University Press.

Ruddick, Sara. 1998. Care as labor and relationship. In *Norms and values: Essays on the work of Virginia Held*, ed. Mark S. Halfon and Joram C. Haber Lanham. Washington: Rowman & Littlefield publishers.

Sangiovanni, Andrea. 2016. How practices matter. *Journal of Political Philosophy* 24 (1): 3–23.

Scanlon, T. M. 2003. Thickness and theory. *Journal of Philosophy* 100 (6): 276–277.

Scheffler, Samuel, 1987. Morality through thick and thin: A critical notice of ethics and the limits of philosophy [Williams 1985]. *Philosophical Review* 96 (3): 411–434.

Sigrid, Mehring. 2011. Medical war crimes. In *Max Planck yearbook of United Nations law*, ed. Armin Bogdandy and Rüdiger Wolfrum, 229–279. Netherlands: Martinus Nijhoff Publishers.

Slote, Michael. 2001. *Morals from motives*. Oxford: Oxford University Press.

Slote, Michael. 2007. *The ethics of care and empathy*. London: Routledge.

Stark, Cynthia. 2010. Abstraction and justification in moral theory. *Hypatia* 25 (4): 825–833.

Steyl, Steven. 2019. The virtue of care. *Hypatia* 34 (3): 507–526.

Steyl, Steven. 2020. Caring actions. *Hypatia* 35 (2): 279–297.

Steyl, Steven. 2021. A care ethical theory of right action. *The Philosophical Quarterly* 71 (3): 502–523.

Tronto, Joan. 1993. *Moral boundaries: A political argument for an ethic of care*. New York: Routledge.

Tronto, Joan. 2010. Creating caring institutions: Politics, plurality, and purpose. *Ethics and Social Welfare* 4 (2): 158–171.

Tronto, Joan. 2013. *Caring democracy: Markets, equality and justice*. New York: New York University Press.

Vayrynen, Pekka. 2011. Thick concepts and variability. *Philosophers' Imprint* 11 (1): 1–17.

Williams, Bernard. 1985. *Ethics and the limits of philosophy*. Cambridge, MA: Harvard University Press.

Zangwill, Nick. 1995. The beautiful, the dainty and the dumpy. *British Journal of Aesthetics* 35 (4): 317–329.