Document generated on 09/15/2026 1:49 p.m.

Canadian Journal of Bioethics
Revue canadienne de bioéthique

Canadian Journal of Bioethics
Revue Canadienne de Bioéthique

Getting Real About Killing and Allowing to Die: A Critical Discussion of the Literature

Andrew Stumpf and Dominic Rogalski

Volume 4, Number 2, 2021

Éthique et AMM
MAiD Ethics

URI: https://id.erudit.org/iderudit/1084448ar
DOI: https://doi.org/10.7202/1084448ar

See table of contents

Publisher(s)

Programmes de bioéthique, École de santé publique de l'Université de Montréal

ISSN

2561-4665 (digital)

Explore this journal

Cite this article

Stumpf, A. & Rogalski, D. (2021). Getting Real About Killing and Allowing to Die: A Critical Discussion of the Literature. *Canadian Journal of Bioethics / Revue canadienne de bioéthique*, *4*(2), 8–24. https://doi.org/10.7202/1084448ar

Article abstract

The moral significance of the distinction between killing and allowing to die has played a key role in debates about euthanasia and physician assisted suicide. Since the withdrawal of life-sustaining treatment is held as morally permissible in the medical community, it follows that if there is no morally significant difference between killing and allowing to die, then there is no morally significant difference between withdrawing life-sustaining treatment or administering a lethal injection to end a patient's life. Consistency then requires that voluntary active euthanasia (VAE) is also morally permissible. The debates over whether the distinction is morally significant have carried on for decades with little hope of consensus. We begin by surveying the literature to identify common argumentative strategies used in defending or rejecting the distinction's significance. We observe, based on our review, that many of these strategies operate in ways that are conceptually removed from the concrete clinical situation of physicians involved in practices that lead to patient death (by withdrawal of treatment or VAE). We conclude by arguing for a novel way of moving the debate forward indicated by our reading of the literature, namely, by paying careful attention to the moral experience of physicians involved in end-of-life interventions to understand how they experience these practices. Exploring physician experience can reveal how the distinction may or may not be useful for moral deliberation and can provide the needed context to theorize about the distinction in a more empirically informed and practically useful way.

© Andrew Stumpf and Dominic Rogalski, 2021

This document is protected by copyright law. Use of the services of Érudit (including reproduction) is subject to its terms and conditions, which can be viewed online.

https://apropos.erudit.org/en/users/policy-on-use/

érudit

This article is disseminated and preserved by Érudit.

Érudit is a non-profit inter-university consortium of the Université de Montréal, Université Laval, and the Université du Québec à Montréal. Its mission is to promote and disseminate research.

https://www.erudit.org/en/

Canadian Journal of Bioethics
Revue Canadienne de Bioéthique

A Stumpf, D Rogalski. *Can J Bioeth / Rev Can Bioeth.* 2021;4(2):8-24

ARTICLE (ÉVALUÉ PAR LES PAIRS / PEER-REVIEWED)

Getting Real About Killing and Allowing to Die: A Critical Discussion of the Literature

Andrew Stumpf[a], Dominic Rogalski[b]

Résumé
La signification morale de la distinction entre tuer et laisser mourir a joué un rôle clé dans les débats sur l'euthanasie et le suicide assisté. L'arrêt d'un traitement de survie étant considéré comme moralement acceptable par la communauté médicale, il s'ensuit que s'il n'y a pas de différence moralement significative entre tuer et laisser mourir, il n'y a pas non plus de différence moralement significative entre l'arrêt d'un traitement de survie et l'administration d'une injection létale pour mettre fin à la vie d'un patient. La cohérence exige donc que l'euthanasie active volontaire (EAV) soit également moralement admissible. Les débats sur la question de savoir si cette distinction est moralement significative ou non se poursuivent depuis des décennies, sans grand espoir de consensus. Nous commençons par passer en revue la littérature afin d'identifier les stratégies argumentatives communes utilisées pour défendre ou rejeter l'importance de la distinction. Nous observons, sur la base de notre revue, que nombre de ces stratégies opèrent d'une manière qui est conceptuellement éloignée de la situation clinique concrète des médecins impliqués dans des pratiques qui conduisent à la mort du patient (par arrêt de traitement ou EAV). Nous concluons en plaidant pour une nouvelle manière de faire avancer le débat indiquée par notre lecture de la littérature, à savoir en accordant une attention particulière à l'expérience morale des médecins impliqués dans les interventions de fin de vie afin de comprendre comment ils vivent ces pratiques. L'exploration de l'expérience des médecins peut révéler comment la distinction peut ou non être utile à la délibération morale, et peut fournir le contexte nécessaire pour théoriser sur la distinction d'une manière plus empirique et plus utile en pratique.

Mots-clés
éthique clinique, bioéthique, tuer et laisser mourir, faire et permettre, causalité, intention, abstraction, contexte clinique

Abstract
The moral significance of the distinction between killing and allowing to die has played a key role in debates about euthanasia and physician assisted suicide. Since the withdrawal of life-sustaining treatment is held as morally permissible in the medical community, it follows that if there is no morally significant difference between killing and allowing to die, then there is no morally significant difference between withdrawing life-sustaining treatment or administering a lethal injection to end a patient's life. Consistency then requires that voluntary active euthanasia (VAE) is also morally permissible. The debates over whether the distinction is morally significant have carried on for decades with little hope of consensus. We begin by surveying the literature to identify common argumentative strategies used in defending or rejecting the distinction's significance. We observe, based on our review, that many of these strategies operate in ways that are conceptually removed from the concrete clinical situation of physicians involved in practices that lead to patient death (by withdrawal of treatment or VAE). We conclude by arguing for a novel way of moving the debate forward indicated by our reading of the literature, namely, by paying careful attention to the moral experience of physicians involved in end-of-life interventions to understand how they experience these practices. Exploring physician experience can reveal how the distinction may or may not be useful for moral deliberation and can provide the needed context to theorize about the distinction in a more empirically informed and practically useful way.

Keywords
clinical ethics, bioethics, killing and letting die, doing and allowing, causation, intention, abstraction, clinical context

Affiliations
[a] Department of Philosophy, St. Jerome's University, Waterloo, Canada
[b] Faculty of Science, University of Waterloo, Waterloo, Canada
Correspondance / Correspondence: Name, adhstumpf@uwaterloo.ca

INTRODUCTION

One popular argument for the moral justifiability of voluntary active euthanasia (VAE) goes like this:

1. Allowing a patient to die by withdrawing life-sustaining treatment (WLT) is uncontroversially morally acceptable.
2. There is no morally significant difference between allowing a patient to die (by withdrawing treatment) and killing a patient (by VAE).[1]

Therefore (if we are being consistent),

3. VAE must be regarded as morally acceptable.

[1] In the Canadian context, medical assistance in dying (MAID) includes both what has traditionally been called physician assisted suicide (PAS) and voluntary active euthanasia (VAE). In this paper, we focus on the contrast between VAE and allowing to die by withdrawal of treatment, but *mutatis mutandis*, our approach could be extended to PAS and to other forms of allowing to die such as refusal to initiate life-sustaining treatment or ("terminal") sedation proportionate to the relief of pain and suffering.

2021 Stumpf & Rogalski. Creative Commons Attribution 4.0 International License
ISSN 2561-4665

In the debate over this argument, the acceptability of the second premise is the focus, since there is broad if not complete consensus concerning the first premise, and together with the first, the second premise gets us to the conclusion in a tidy and effective way.[2] So it is unsurprising that an industry has grown around supporting and attacking premise two. Over the last four decades at least, the debate has raged on with few signs of consensus emerging (1).[3]

We begin this paper by surveying the literature to identify common argumentative strategies used in defending or rejecting premise two.[4] Based on our review, we find that many of these strategies operate in ways that are conceptually removed from the concrete clinical situation of physicians involved in bringing about patient death by WLT or by VAE (administering lethal injections). Consequently, debate in this area has been bogged down by a tendency toward irrelevant abstractions. We argue instead for a novel way of moving the debate forward, namely, by paying careful attention to physicians' sense of their own moral agency and moral responsibility, in order to find out whether physicians involved in end-of-life interventions experience themselves to be doing basically[5] the same thing in WLT and VAE. If physicians experience these two practices as *different*, and if they articulate the difference in terms used in the bioethics debate, that would suggest that the bioethical distinction between killing and allowing to die makes a difference for practical moral deliberation and action in the clinical setting. If, however, physicians experience the two practices as equivalent despite an awareness of the way the distinction can be drawn, that would suggest that the bioethical distinction *does not* make a difference for practical deliberation and action in the clinical setting. Given that the usefulness of a distinction for real moral agents engaged in relevant contexts of deliberation partly constitutes that distinction's moral significance, greater attention to physicians' experience should help move us closer to discerning the significance, or lack thereof, of the distinction between WLT and VAE.[6]

ANALYZING THE DISTINCTION AND SAYING WHETHER IT MATTERS

We begin our review of the debate by describing the main accounts of the distinction between killing and allowing to die. We then delineate the common argumentative strategies authors have deployed to support or reject the claim that the distinction between killing and allowing to die is morally significant. The question of what the distinction consists in is a question for the metaphysics of action. Attempting to answer the further question of how the (metaphysical) distinction matters for morality brings us into the domain of normativity. The present section explores four common attempts to account for the distinction, in terms of 1) *doing vs. allowing*, 2) *difference in causal explanatory roles*, 3) *intention vs. acceptance* and 4) the *broader narrative context*. In each case we describe the accounts and show how their advocates attempt to use them to establish the distinction's moral significance and how their critics resist these efforts.[7]

The Doing vs. Allowing Account

According to the first way of explaining the distinction, killing should be understood as an active *doing* or commission that differs from (merely) *allowing* to die. One may allow another to die in two ways: either by one's inaction (omission) or by removing an obstacle that would have prevented the outcome (death) from occurring (2-4).[8] Administering a lethal injection, on this account, is a positive act of killing, while refusing to initiate life-sustaining treatment (an omission), or withdrawing life-sustaining treatment (removing an obstacle preventing death) is not killing, but only allowing to die.[9] And since morally wrong commissions are harder to justify (and are often held to be morally worse) than morally wrong cases of allowing, the distinction between active doing and allowing (by omission or by removing an obstacle) makes a moral difference.

Critics of the Doing/Allowing Account have been quick to point out two objections. First, it seems quite possible to kill by omission. This can be seen in the case of the Enemy or Interloper, where someone (say, a greedy nephew) pulls the plug on another's life-support in order to hasten death (and, say, secure an inheritance more quickly) (5,6).[10] Like the Doctor who withdraws life support, the Enemy removes an obstacle preventing death and so allows the patient to die. Our intuitions tell us that the Enemy has killed the patient, but the Doing/Allowing account does not (at least, not on its own), support that

[2] There are individuals who strongly oppose withdrawing life support from patients, even when the patient is wholly brain-dead. Yet, mainstream North American legal and medical thought allows for WLT. Thanks to a reviewer of this paper for reminding us that it would be misleading to claim that premise 1 is entirely uncontroversial.

[3] Chandler, writing in 1990, notes a serious lack of consensus in this debate, the start of which he traces to Bennett (1966). Thirty years after Chandler's comment, things have not changed substantially.

[4] Our survey involved searching the following databases for sources that matched our inclusion criteria: PubMed, Ovid Medline, Scopus, Web of Science, CINAHL, Philosopher's Index, PhilPapers, JSTOR, and POIESIS. Inclusion criteria broadly consisted of whether resources were directly dealing with Killing/Allowing to Die (and related distinctions), whether there was a clear bioethical argument, and the citation frequency or reference within other sources.

[5] Clearly physicians are not doing *identically* the same thing with WLT and VAE. The question is whether, given their relationship to the patient and the nature and purpose of the interventions, physicians experience the two procedures as functionally and morally equivalent.

[6] The reader might worry that focusing on physician experience is an unwarranted privileging of the physician's moral perspective over the perspectives of others involved in end-of-life decisions and practices (i.e., patients, patient advocates, ethics panels, clinical ethicists, lawyers, judges, nurses, and social workers). A full treatment of the concrete clinical situation should indeed consider each of these voices. To make a manageable beginning, we have selected one centrally important observational angle. Our thanks to Nick Ray for bringing this point to our attention.

[7] These four accounts can overlap in different ways and should not be understood as fully independent, mutually exclusive, or exhaustive.

[8] McGrath refers to the second type of allowing as "causation by disconnection" (2, p.91). Bronner, against Miller and Truog (28), follows Foot's account of removing a barrier as "merely allowing a pre-existing threat to cause harm"; in the case of WLT, no new sequence of events has been set in motion, but a barrier that was holding back a pre-existing sequence has been removed (4).

[9] Many speak of WLT as an instance of an omission. Speaking precisely, when it leads to death, WLT seems to have the form of allowing to die by removing an obstacle that was holding back death. Describing WLT as an omission (omitting to provide further treatment) fails to capture the fact that treatment was provided in the first place and is now being discontinued. Nevertheless, we will not quibble in what follows with those who classify WLT as allowing to die in the sense of omitting to do something to prevent death.

[10] We discuss these sorts of comparisons in later sections of the paper.

judgment or distinguish the Enemy from the Doctor. Second, critics object that certain allowings can be just as bad as corresponding commissions, so that even if we can clearly distinguish between acts of killing and allowing to die, this distinction does not matter for morality. The contrast strategy of Rachels, Lichtenberg, Tooley, and others (7-10) is an – by now famous – instance of this type of objection. By describing paired cases judged to be morally equivalent but that differ only in that one is a killing and the other an allowing to die, the contrast strategy purports to show that doing versus allowing does not make a moral difference. Others argue that the distinction does not matter morally because it is the *behaviour* and *outcomes* that are truly important in our moral assessments, not the nature of the act itself (11-13).[11]

Defenders of the Doing/Allowing account have begun to recognize that the mere distinction between positive acts and allowings does not suffice to pinpoint the morally significant difference between killing and allowing to die. Merkel (14), for instance, argues for a "normative asymmetry" between omissions and commissions grounded not in the nature of doing and allowing themselves, but in the fact that different obligations are relevant to doings and to allowings. This asymmetry is reflected in law[12]:

> omissions constitute legal liability for ensuing consequences (e.g. death) only if the omitting person had a legal obligation to act and thereby prevent the particular consequence. Such an obligation presupposes a special legal duty owed especially to the beneficiary. In contrast, the obligation to refrain from actively killing another is a general, basic and permanent duty of everyone towards everyone else (certain narrowly defined exceptions like self-defense notwithstanding). (14, p.353)

Merkel notes that one is culpable in a different way for pushing a child into a pond and thereby causing their death, than for refraining from assisting a child who is drowning. One who carried out the former act would be culpable under the general prohibition against murder. For the latter inaction, one would be culpable only if one carried a special duty to preserve the life of the victim, for instance if the victim was one's own child. McLachlan (15) also grounds the significance of the distinction between omission and commission on the differing moral duties that apply. The general obligation to refrain from killing other persons differs from the general obligation to try (or continue to try) to prevent people from dying. So, even if we have the latter obligation (which is questionable at best), someone who is guilty of both would be guilty of violating two distinct obligations (15).[13] An important objection to such efforts argues that on careful analysis of the distinction between doing and allowing, most morally interesting instances of doing can also be described as cases of allowing (2).[14] If doing and allowing do overlap, the doing/allowing distinction will not explain the moral difference between, for instance, the case of the Doctor who performs WLT and the Enemy who disconnects life support to get an inheritance more quickly. More to the point, on McGrath's (2) analysis, a doctor's act of switching off a patient's life support at the patient's request must be characterized both as a doing and an allowing.[15] But this undercuts the attempt to distinguish WLT from VAE on the grounds that the former is an allowing and not a positive doing.

Merkel and McLachlan contrast WLT (considered as a type of allowing to die) with VAE (considered as active killing), and VAE comes out morally worse because killing always violates a very strong moral obligation while allowing to die either does not violate any obligation or (when wrongful) violates different, weaker obligations. Woollard (16, p.464) also argues that the distinction between doing and allowing systematically "links up with more fundamental moral concepts," such as the authority necessary for ownership of one's own body and mind (17,18). But instead of contrasting the practices of WLT and VAE themselves, Woollard contrasts two kinds of refusals related to the two practices – the refusal of requested WLT and the refusal of requested VAE. The idea is that requests for WLT impose obligations in a way that requests for VAE do not. Refusing a request for WLT and continuing to administer unwanted life-sustaining treatment involves actively doing harm and a violation of strong rights. Refusing a request for VAE, on the other hand, is morally less problematic, because it merely allows harm and does not violate any strong rights. According to Woollard, doing harm is generally harder to justify than allowing harm, because doing harm involves imposing on the victim, while allowing harm involves refusing to be imposed upon by the victim. This difference between doing and allowing is a necessary entailment of the claim that anything, including our own bodies, belongs to us. It offers us "*prima facie* protection against the behavior or needs of others intruding upon what belongs to us" (16, p.465).[16]

---

[11] As we will see in the fourth account of the distinction, some argue that discussion of the distinction can only be meaningful within a larger conversation about whether a utilitarian approach to ethics is preferable to a deontological approach, or vice versa. See Chandler (1).

[12] As Merkel notes, while the domain of the law is distinct from that of ethics, legal obligations often reflect ethical ones. To the extent that they do so, an analysis of the obligations from a legal point of view can inform ethical analysis.

[13] Paterson makes a similar point: A judgment that an act and an omission are morally equivalent or not "can only flow from an analysis of prior duties, not independently of them" (27, p.31). Specifically, an omission that brings about death as an outcome can be just as morally bad as an active commission (killing) in the event of a strong duty to act (to preserve life) that has been violated by the omission. But given that killing always violates the (very strong) moral duty not to kill a person, while allowing to die can and frequently does occur in the absence of any strong duty to act to preserve life, it is true that there is at least a *prima facie* (metaphysical) difference between killing and allowing to die that, in connection with applicable moral duties, makes a difference from a moral point of view.

[14] McGrath (among others) prefers the term "making" to "doing." We do not believe anything significant turns on this terminology, so for the sake of consistency we use "doing" in characterizing McGrath's position.

[15] Defenders of the Doing/Allowing Account may wish to question McGrath's understanding of doing something (making something happen) as *causing that something to occur by connection*, where "causing by connection" is illustrated by paradigmatic cases like one billiard ball striking another and thereby making it move (2, p.91). In our discussion of Causation accounts in the next section, Bronner (4) provides some reason to think that the idea of causation relevant to WLT involves something quite different. Furthermore, our analysis in the second part of the paper will suggest that the real problem here is the attempt to classify WLT using abstract notions of doing and allowing that ignore concrete contextual features of the situation in which WLT standardly takes place. The Enemy (unlike the Doctor who withdraws treatment legitimately) has no duty to remove treatment or any relational or institutional authorization to intervene at all. In the end, one may not be able to account clearly for the difference between the Enemy and the Doctor without considering their respective intentions.

[16] As Woollard notes, the protection is not absolute; sometimes "considerations of harm and benefit can outweigh the authority of possession" (16, p.465).

Woollard's account distinguishes VAE from WLT on the grounds that failure to comply with a patient's request to withdraw treatment would constitute an unjustified doing of harm, while failure to comply with a request for VAE would not involve doing (but only allowing) harm to the patient. Others make the related point that when a patient requests to be allowed to die (by WLT), the patient's authorization involves the refusal of an intervention, and must be granted based on traditional rights of non-interference. By contrast, the request to be killed (by VAE, for instance) is a request *for* an intervention, which, if required at all, would be grounded in more tenuous entitlement rights. Law reflects the judgment that there is a much greater obligation to honour refusals than to grant requests (19). As Wolf puts it, "The negative right to be free of unwanted bodily invasion… is clearly grounded in fundamental entitlements to liberty, bodily privacy, and freedom from unconsented touching; in contrast there is no clear 'right' to kill yourself or be killed" (20, p.300; 21). Opponents to such moral and legal claim-right arguments turn the question around by observing that moral and legal obligations do not *prohibit* requests for VAE. Miller et al. (22, p.458) state that "patients arguably have a moral liberty-right of non-interference by others with their physician's voluntary compliance with a valid request for… active euthanasia, not a claim-right to receive… [a] lethal injection upon demand from an objecting and unwilling physician." These authors claim that greater focus should be given to determining the circumstances in which a physician's compliance to competent patient's voluntary request for death is legitimate and where such compliance and requests are not legitimate (22,23). If Miller et al. are correct to focus on the negative right to non-interference and the circumstances in which requests for VAE are legitimate and where they are not, then it suggests that the moral and legal claim-right arguments of Woollard and others are more suited for addressing positions of conscientious objectors or those who claim that such requests and compliance are never legitimate (i.e., the circumstances are such that only allowings are sometimes permissible).

In both its forms, the Doing/Allowing account as applied to VAE and WLT is held to matter for morality because harmful active commissions (doings) are generally harder to justify than harmful instances of allowing to die. On the one account (Merkel, McLachlan) there is a higher moral justificatory threshold for killing than for allowing to die. On the other (Woollard, Wolf), doing harm by refusing requests for WLT are harder to justify than refusing requests for VAE and thereby allowing harm to occur. Disagreement over the Doing/Allowing account remains over whether doing and allowing are mutually exclusive, and over whether the focus should be on the purported positive right to receive VAE or the negative right to non-interference in receiving VAE.

The Causation Account

Like Doing/Allowing Accounts, Causation Accounts attempt to characterize the difference between killing, as actively *doing* something, and merely *allowing* an outcome (death) to happen. But Causation Accounts unpack this difference in terms of the different roles played by killing and allowing to die in causal explanations. Accordingly, when I kill someone, I must be understood as causing their death, but when I merely allow someone to die, it is not appropriate to describe myself as having caused their death (24,25).[17] More specifically, in killing a patient by administering a lethal injection, a doctor assumes (causal) responsibility for the patient's death, but it would be incorrect to describe a doctor's merely allowing their patient to die (by WLT) as causally responsible for the death, since it is the underlying medical condition that (proximately) causes death (25-27).[18] Stauch (25), for instance, argues for a necessary link between an act (commission) and causal authorship for its outcome which is not present in omissions that bring about the same outcome when no socially imposed duty to act obtains.[19] Accordingly, a doctor who removes life-sustaining treatment from their patient is only responsible for the patient's death when causal status can be ascribed to their omission, and such causal status is relevant only in the context of a socially imposed duty to treat. By contrast, a doctor who engages in VAE necessarily assumes authorship over the patient's death and their act necessarily possesses causal status.[20] The Causation Account combines such reasoning with the assumption that it is always wrong for a health-care professional to assume causal responsibility for a patient's death, to claim that the distinction between killing and allowing to die makes a moral difference.

The above characterization of the Causation Account is often referred to as "the standard view" in traditional medical ethics and law (28). The arguments collectively advanced by Miller, Truog, and Brock (22,28) are aimed against this view. In general, their argument is that, understood correctly, WLT is a species of 'killing' rather than an 'allowing to die.' They contend that classifying WLT as a case of allowing to die springs from an attempt to avoid conflict with the traditional moral

[17] In this paper we avoid a full discussion of the causal status of omissions (how could an *inaction* – a *not*-acting – "cause" anything?) and assume that there is a perfectly good and usable sense of "cause" in which it is obvious that we can cause outcomes by our inactions. This can be established with cases. Selena is the Prime Minister's bodyguard and does not act to stop an attack on the Prime Minister's life when they could easily have done so. Selena has clearly contributed causally to the Prime Minister's death by their inaction. Similarly, the gardener's job is to water the flowers, but the gardener does not water them when they easily could have. The gardener has caused, by their inaction, the withering of the flowers. Hart and Honoré note that we attribute causal status to omissions when they involve deviation from an established and regularly maintained standard or routine: "The gardener's failure to water the flowers… [unlike the failure of everyone else to water them] stands on a different footing. It is not merely a breach of duty on his part, but also a deviation from a system or routine" (24, p.38; 25).

[18] Here is one of several places where the Doing/Allowing Account and the Causation Account overlap. McLachlan (15, p.637) distinguishes causes (properly speaking) from omissions, classifying a typical omission as the "absence of a preventative anti-cause." While an omission, like a (proper) cause, is often a necessary condition for an event to occur, the omission itself cannot directly kill, but can only (indirectly) lead to death insofar as it fails to prevent it. An omission can still be wrong, but when it is, it is a wrongful omission, rather than a wrongful killing.

[19] Note that the distinction between types of allowings is blurred here; McGrath and others would classify WLT more precisely as the removal of an obstacle holding back death rather than as an omission.

[20] See Walton (13) for a related argument that links the language of action with the language of control. Here, when an agent takes positive action, they take control and interrupt "the course of nature," whereas an omission does not assume control and is less active in bringing about an outcome (13).

norm that doctors must not kill by denying that the conduct causes death.[21] Contra the standard view, Miller and Truog employ a common-sense notion of causation which maintains that "causes are events that make a difference in explaining a particular occurrence" (28, p.6).[22] In the case of WLT, the withdrawal makes a difference in explaining the timing of death (in some cases), thus it contributes to causing death by initiating a fatal sequence (17,24). Pointedly, Miller and Truog state,

> The prevailing view that withdrawing LST [life-sustaining treatment] does not cause death has been propounded as a universal proposition applying to all cases when physicians have no duty to continue treatment. With respect to patients who could live indefinitely with continued LST, it is most patent that withdrawing LST causes death. Recognizing that it causes death in these cases points the way to seeing that it also does so, at least as a contributory cause, in more typical cases, where death is hastened…by the withdrawal of treatment (28, p.4-5).

For Miller et al. (28) WLT can cause death but causing death does not affect its ethical legitimacy.[23] That is, even if it is granted that WLT causes death, it does not follow that this routine clinical practice must be abandoned; instead, it suggests that the problem is with the moral norm that prohibits ever causing death to patients, and that this problematic norm is what should be reconstructed. Once the problematic norm is modified, then the standard moral assessment that treats WLT and VAE as morally distinct does not hold and so it must be rejected.

Proponents of the Causation Account need not deny that there is an obvious sense in which the withdrawal of life-sustaining treatment causes death: were it not for the removal of the ventilator, the patient would still be alive. The withdrawal of the treatment makes up part of the total causal story needed to account for the outcome. Nevertheless, defenders of this account highlight the fact that only certain elements within that total causal story are worth mentioning for the purpose of correctly explaining the outcome. When the patient has been merely allowed to die, they argue that the correct attribution of the proximate cause of death would be the underlying medical condition that *prompted the use of life-sustaining treatment in the first place* (25-27). McGee (29,30) takes this approach to challenge Miller, Truog, and Brock's collective arguments.

McGee defends the standard view that WLT should be classified as allowing to die rather than killing. McGee (30) notes that Miller and Truog (28) maintain that *withholding* LST is an omission and does not cause death, whereas *withdrawal* of LST (WLT) involves acts that have causally significant consequences, thus WLT causes death and is more similar to VAE (killing) than to withholding LST (allowing to die).[24] McGee argues that WLT can be classified as letting die despite acts being involved, as the initiation of LST in the first place is morally significant. It is significant because it 'causes' only by allowing the initial dying process to resume, not by doing or creating a new fatal sequence, as Miller and Truog suggest. Therefore, the effect of the act (WLT) is that physicians stop saving the patient and allow them to die (29-31). This account is used to explain the differing characterization of conduct between the Doctor and the Enemy in WLT (5,6).[25] Physicians initiate LST and they are in a position to validly WLT, whereas others (e.g., greedy nephews) do not initiate LST and are not in a position to WLT. So unauthorized WLT can be consistently classified as killing rather than as letting die, as the identity of moral agents influences how an act is classified and is morally significant (30,31).[26]

Bronner (32) provides an analysis against the standard view in the context of novel medical technologies that is resistant to critiques proposed by McGee and others (33). Focusing on the concrete case of deactivating a total artificial heart (DAH), Bronner argues that the standard view cannot accommodate DAH, even if it can accommodate other forms of WLT.[27] A total artificial heart, Bronner claims, is relevantly similar to Implantable Cardioverter Defibrillators (ICDs), except with respect to deactivation: deactivation of ICDs allows a patient to die from a pre-existing condition, whereas DAH kills since there is no longer the pre-existing lethal condition after the intervention. That is, the causal story, as held by the standard view, is broken since there is no longer an underlying medical condition that (proximately) causes death after implantation of the artificial heart. As such, DAH presents a dilemma for medical ethics; it demonstrates conflict between the permissibility for a physician to discontinue any treatment upon request from a competent patient and the moral norm that it is never permissible to intentionally kill a patient.

Elsewhere, Bronner argues that Brock, Miller and Truog wrongly assume that this "common-sense" notion of causation (causation as difference-making) is relevant to the context of decisions to withdraw life-sustaining treatment (4). In such a

[21] We elaborate on their case for this point in the second part of this paper.
[22] That is, *x* causes *y* to die if and only if *x* makes the difference between *y* being alive or dead. The point is to distinguish relevant causes from the conditions necessary (but not sufficient) for a cause to produce its effect.
[23] Miller and Truog (28) stress that we can attribute causes without holding people culpable for causing death. More generally, they maintain that causal judgments are independent of ascribing moral responsibility, so that causation does not equate to moral culpability.
[24] It might be asked how Miller and Truog can consistently hold that withholding treatment is not also a form of killing, given that withholding can make the difference between a patient's being alive or dead. They appear to distinguish withholding treatment from withdrawing treatment by noting that the former involves omission and thus does not cause death, but the latter involves positive action and must be regarded as the cause of death when it occurs (28, see p.6,13-14).
[25] We considered the case of the Enemy or Interloper as an objection to the Doing/Allowing Account above and will return to it in the second part of this paper.
[26] In Asscher's terminology (34), the physician has taken responsibility for the patient's care (within the context of a system of care with its standards and protocols) and thus stands in a very different relationship to the patient than the Enemy in respect to such an action. In part two, we will consider relational and institutional context as key aspects of the clinical situation relevant to an adequate moral assessment of either WLT or VAE.
[27] Bronner (32) is careful to note that he neither argues that DAH kills while other forms of WLT merely allow to die, nor takes a stand on whether the standard view can accommodate other forms of WLT. The aim is to show a concrete example in which there is conflict in the principles established in the standard view, namely, 1) Discontinuation "it is permissible to discontinue any treatment upon the request of a competent patient" and 2) Prohibition "it is never permissible for a physician to intentionally kill a patient" (32, p.347-348).

context, it is understood that WLT will make the difference in the patient's dying or not. But that is not the issue of concern. Instead, the focus is on the role of the doctor, specifically on whether the doctor's act violates the prohibition against causing the death of their patient. The sort of causation in question, in attempting to answer *that* question, is not causation as difference-making but causation as doing (4). We want to know, in other words, whether or not the doctor's involvement in the patient's death falls on the doing side of the doing/allowing distinction.[28] Understood in this way, the distinction between killing and allowing to die remains morally significant; it contrasts a doctor's *doing* lethal harm to the patient with a doctor merely allowing lethal harm (and, we might add, doing and allowing harm carry different justificatory thresholds and/or relate to different moral obligations). Bronner argues that understanding doing (not difference-making) as the relevant notion of causation in the context of WLT has the effect of undermining many arguments against the view that WLT is an act of letting die rather than killing since those arguments typically assume causation as difference-making.

In our discussion of the Doing/Allowing Account, we noted that omissions are culpable or not depending on what obligations obtain in a given situation. Similarly, whether allowing a person to die should be attributed as a cause of that person's death may depend on the presence or absence of a duty to act (e.g., to provide the person with treatment, or to refrain from stopping ongoing treatment). As Stauch (25) puts it, when a doctor allows a patient to die in the absence of a duty to treat, the doctor does not assume causal authorship over the patient's death. In such a case, the doctor's omission becomes equivalent to similar omissions by everyone else in the world that also did not ensure continuing life-sustaining treatment. No such omissions are morally culpable. By contrast, it is arguably not possible to kill a patient without "assuming causal authorship" in relation to their death (25,34,35).[29] Paterson concurs:

> [T]he strict notion of cause is not capable of grounding the difference between killing and letting die. Physical causes do not explain an agent's accountability unless they are interpreted through the lens of moral responsibility. However, they do form a necessary condition for the evaluation of human conduct, which, taken in conjunction with a moral duty not to intend to kill, can so ground the distinction. (27, p.34)

The acknowledgement that the culpability and causal status of omissions depend on their relations to relevant obligations is an important step away from abstract thinking and toward a greater sensitivity to concrete aspects of the clinical setting. We will consider such positive cases of concrete moral reasoning more fully in the second part of this paper. We conclude our discussion of the Causation Account by noting that there remain difficult questions both about the general nature of causation and about the type of causation that should be considered in comparing WLT and VAE.

The Intending vs. Accepting Account

A third way of accounting for the distinction between killing and allowing to die appeals to the way the will is involved in different ways of acting. Killing, it is argued, necessarily involves intending someone's death, while allowing to die does not necessarily involve such an intention, even when the death is foreseen and accepted. Since acting to intentionally end an innocent person's life violates a basic moral prohibition, killing is always wrong, whereas allowing to die need not be wrong, though it may be (19,36).[30] Those who want to account for the distinction by claiming that it is possible to foresee and accept but not intend the patient's death in WLT must respond to those who reject such efforts as artificial and practically useless at best, and absurd at worst (11,27,28,37,38). Many such rejections of intention accounts focus their efforts against the doctrine of double effect. The doctrine is invoked to explain or justify an action that has both a good and a bad or harmful effect (38). In WLT, the good (intended) effect is respecting patient autonomy by removing burdensome treatment and the bad (unintended) effect is causing or permitting death. Conversely, in VAE, the good effect is asserted to be respecting patient autonomy by treating intolerable suffering and the bad effect is causing death. Opponents of VAE argue that it produces the good effect by intending the bad effect as a means – the patient's death is intended as a means to ending their suffering in accordance with their wishes. By contrast, WLT does not necessarily involve intending death as does VAE.

Miller and Truog criticize double effect by pointing out that it begs the question to characterize death as a bad or harmful effect in the context of end-of-life decisions "unless death is viewed (unreasonably) as an objective harm under all circumstances" (28, p.16). The doctrine itself does not tell us what values are important or correct for a patient at the end of life; and continued existence may be a greater harm from the patient's perspective than death by VAE. If a patient's values and preferences matter in the context of end-of-life decisions, then the traditional prohibition against causing death, as typically invoked when appealing to the doctrine of double-effect, may be an untenable assumption. According to a second line of objection against double-effect, agents are morally responsible for more than what they directly intend, such that if one does something with avoidable and foreseeable consequences, then those consequences have been intended (11,39-42). Since in WLT the patient's death is avoidable and foreseeable, the patient's death should be considered to have been

[28] Bronner (4) assumes that the difference between doing and allowing is sufficiently clear on the basis of intuitions, stimulated by reflection on cases.

[29] Stauch's phrase, "assuming causal authorship" for an outcome (such as death) is strikingly similar to Asscher's notion of "taking responsibility" (34, p.282). Asscher argues that killing is morally distinct from letting die because killing involves taking responsibility for a person's death while allowing to die does not involve taking such responsibility. Asscher qualifies this general statement by noting that if the agent is already responsible for the situation, the distinction is no longer morally relevant, and allowing a person to die in such a circumstance could be as morally bad as (or even worse than) killing. See Trammell (35) for a partial precursor of Asscher's Responsibility Account.

[30] Contra the Causation Account, Bishop (36) argues that since both acts of omission and acts of commission can result directly from human intention, they can both be classified as direct acts of will and, if formed by the intention to kill, should both be viewed as forms of directly causing death. Bishop distinguishes from both of these the case of allowing to die of natural causes in which one does not intend the death of the patient but merely allows it for the sake of some higher good. In this case, Bishop argues, death should be seen as an indirect result of the action and a direct result of the underlying disease.

intended in WLT as much as in VAE. A third criticism asserts that empirical evidence shows that clinicians often intend patient death in cases of WLT (28,43,44). Miller and Truog (28) argue that these empirical observations suggest a problem with the moral norm that universally prohibits causing patient death, not with the routine and legitimate practice of WLT.[31]

Sulmasy defends the coherence of the Intending vs. Accepting Account, provided that one understands the context of usage of those who have wished to defend the "traditional view" that killing is morally distinct from allowing to die (19). Those who have wished to distinguish WLT from VAE by claiming that the former was an instance of allowing to die and the latter an instance of killing, he argues, have understood these terms in a particular way. Specifically, they have understood "killing" to mean "an act in which an agent creates a new, lethal pathophysiological state with the specific intention in acting of thereby causing a person's death," and "allowing to die" to mean "an act in which an agent either performs an action to remove an intervention that forestalls or ameliorates a pre-existing fatal condition or refrains from action that would forestall or ameliorate a pre-existing fatal condition, either with the specific intention of acting that this person should die by way of that act or not so intending" (19, p.57-58). Once we understand this usage, the moral relevance of the distinction, based on the different intentions involved in WLT and VAE becomes clear, provided that we also have a correct understanding of the intentions. Sulmasy's account of intentions appeals to the notion of "commitment": a physician can foresee and even desire that a patient should die quickly after WLT and yet not intend the patient's death because they are not committed to the death as the condition that fulfills their intention (19, p.59).[32] Whether or not an agent is committed to an outcome (e.g., death) in this way can be tested by asking such questions as, "If the patient's death did not take place, would you consider your action to have failed? Would you try some other way to bring about their death?" If the answer to these questions is "yes", then the agent is committed to the outcome; it is what their action was *aiming at*, what their action *intended*. If, on the other hand, the agent can answer the questions with a genuine "no", then they did not intend the death by their action; they may have intended instead only that unwanted treatment no longer be given (19,45).

Paterson (27), inspired by Cavanaugh (46-48), provides further analysis of the difference between intending an outcome and merely accepting it by describing intention in terms of "the anatomy of the will":

> To intend evil as an end is to embrace it most deeply within our person. To intend evil means is still to closely embrace evil within the disposition of the will, since means are the necessary vehicles chosen in order to achieve ends. It is only in permitting effects that lie outside the immediate trajectory of the will relating to a choice of end or means that it becomes possible to accept any negative results that may ensure. Here, there is sufficient distance between the direction of the will and resulting effects to allow us to justify the causation of evil (27, p.35).

Accordingly, in carrying out VAE, one necessarily commits oneself to, and thus identifies one's will with, the evil of the patient's death (49).[33] In WLT, one does not commit oneself to the patient's death in this sense; to withdraw treatment, one need only commit to the outcome in which the patient is no longer receiving the treatment.

Because Sulmasy's (19) notion of "intention" appears to differ from that employed in criticisms of the Intending vs. Accepting Account, it is unclear whether those criticisms undermine Sulmasy's version of this account. Miller and Truog (28) and Brody (43) operationalize the term "intention" along the lines of "motivation," and "desire," but not "commitment." The empirical evidence they cite reflects a similar understanding (44,50). Furthermore, Sulmasy's account builds intention into the very act-descriptions used to identify acts as killings or allowings to die. So, in every act of "killing" the agent will intend (be committed to) the patient's death, while the intentions involved in "allowing to die" can vary. This allows Sulmasy to resist arguments from, for instance, Isaacs (51) claiming that the distinction between killing and allowing to die only *seems* to be morally significant because of the way other factors (besides being-an-instance-of-killing or being-an-instance-of-allowing-to-die) come into play.[34]

In summary, the Intending vs. Accepting Account holds that killing always involves the (prohibited) intention that someone will die as a result of my action, while allowing someone to die does not always involve the (prohibited) intention that someone will die as a result of my action. For this reason, there is a morally significant difference between killing and allowing to die. Disagreement remains over how best to characterize the nature of intention, and over whether it is

[31] One might question why Miller and Truog do not conclude merely that many physicians do not, in practice, respect the prohibition on killing patients, rather than point to a problem with that prohibition. In short, they believe that their claims regarding intention hold *when their causation account is accepted*. That is, WLT (and causing patient death) retains its ethical legitimacy irrespective of clinician's intentions (when conducted under the appropriate circumstances, such as with valid consent from a competent patient). They argue that denying causation and holding intentions as fundamentally important to the ethical legitimacy of WLT are driven by commitment to the antiquated moral norm. We return to Miller and Truog's approach in the second section of the paper.

[32] In response to the second line of objection noted above, Sulmasy's account offers an explanation of how it is possible to know (even with certainty or near certainty) that one's action will produce an effect and yet not intend that effect.

[33] Biggar argues that beyond the effect of our actions on the world outside ourselves we must also consider the "subjective, reflexive impact on the agent… of his will's commitment, through intention, to a right or wrong act" (49, p.67). To intend something that is morally evil, on Biggar's view, involves identifying oneself with that evil and thereby corrupting oneself. One effect of such self-identification with evil is the increased likelihood that one will choose to do more of the same sort of evil in the future.

[34] If, as Isaacs sees it, intentions are not intrinsic features of actions, then they may be classified as external factors that sometimes make the difference in our moral evaluations of actions. But if Sulmasy is right to include intentions as integral parts of the actions themselves, then the distinction between killing and allowing to die is itself morally significant. Since Isaacs accepts that there may be a morally relevant distinction between WLT and VAE on the grounds of differing intentions (see p.364), the disagreement between Isaacs and Sulmasy may be merely verbal.

unreasonable to hold that intending death (as an end or a means) is always wrong and, relatedly, whether it is unreasonable to see death as an absolute harm in all circumstances.

The Broader Narrative Context

For the sake of completeness, we briefly consider here some accounts that do not fit easily within the other three main categorizations. Each of these accounts claims that the distinction between killing and allowing to die has moral significance within some larger frameworks of meaning but not in others. Some authors claim that the distinction can no longer be understood today, given the jettisoning, by the dominant culture, of a certain broad narrative context that includes a distinctively Christian attitude toward death (52,53). Meilander (52) argues that VAE but not WLT is incompatible with a conception of care corresponding to this Christian narrative context and attitude. Bishop argues that the distinction between killing and letting die coheres within a traditional "metaphysics of purpose" which links material effects with formal and final causes (intention and will), but fails to cohere within the modern "metaphysics of efficiency," according to which the moral assessment of an act focuses on the effects of the action and their utility (36,45,54).[35] Relatedly, Chandler (1) argues that discussion of the moral significance of killing and allowing to die can only properly take place when couched explicitly in a broader conversation about which ethical theory is correct. Utilitarianism entails that there is no intrinsic moral difference between doing and allowing in general, or between killing and letting die specifically. Deontological ethics and natural law ethics, on the other hand, imply that some allowings are culpable and others are not. If, therefore, our views of the significance of the distinction are a product of the ethical theory we ascribe to, then our intuitions about cases "will merely chart these deeper disagreements" (1, p.420; 55).[36]

In the first part of this paper, we have illustrated each of the three main accounts of the distinction between killing and allowing to die and of its potential moral significance (the Doing/Allowing, Causation, and Intending vs. Accepting accounts). We have also identified the main areas where disagreement remains between advocates of these accounts and those who reject the accounts themselves (at the level of metaphysics of action) or deny that they matter for morality. In addition, we briefly indicated perspectives that seek to push the debate back or up a level by considering broader frameworks of meaning that affect our view of the distinction. Because our aim in this paper is to point out problematic tendencies in the debate and to suggest a new way forward, we do not have space to go into these accounts in more detail. We trust that the preceding discussion is sufficient to acquaint those unfamiliar with, and remind those already familiar with it, of the main contours of the debate up to the present time.

DIAGNOSING THE PROBLEM: IRRELEVANT ABSTRACTION

As we have seen, the debate over the moral significance of the distinction between killing and letting die is extremely complex. Not only is the distinction's significance contested, but the correct way to understand the distinction itself has been questioned from numerous angles. The goal of the present part of this paper is not to resolve the debate, either as a whole or in its parts. Rather, we seek to articulate our uneasiness with some of the approaches commonly taken by participants in the debate, and our sense of the relatively more satisfying nature of other approaches. In brief, we find that treatments of the distinction's significance that are grounded in cases and conceptual analyses removed from the concrete clinical situation tend to be less effective than approaches that attend carefully to important features of the concrete clinical situation. Appreciating what it is that makes some approaches less useful than others will help guide the direction that further debate on these topics should take. We argue that if such debate is to be fruitful, it must keep an attentive eye on relevant features of the clinical situation. Finally, we sketch a proposal for a particular contribution to the debate along such lines, by attending to the lived moral experience of physicians engaged in VAE and WLT.

We begin by demonstrating, using representative examples from the literature, how and why certain approaches to the contested distinction suffer from irrelevance due to abstraction and thereby hinder fruitful debate. We then indicate what makes approaches that pay careful attention to concrete aspects of the clinical situation more helpful and more likely to promote fruitful debate. But first, a brief word to clarify what we mean by "abstraction," "concreteness," "relevance" and "irrelevance" in this context. What is meant by asking, in the clinical context, whether there is a moral difference between killing someone and letting them die? There are in fact three different questions, of increasing specificity, that we might seek to answer here:

1. Is there a moral difference between (actively) doing something that produces an outcome and merely allowing that outcome to occur?
2. Is there a moral difference between causing someone's death (making them die) and merely allowing them to die?
3. Is there a moral difference between stopping someone's heart by administering a lethal injection (at their request) and disconnecting someone from unwanted life-support (at their request)?

[35] Bishop (36) claims that the thinking of Miller et al. falls within this paradigm of the metaphysics of efficiency. See Bishop's *The Anticipatory Corpse: Medicine, Power, and the Care of the Dying* (54) for a much fuller and more historically and methodologically grounded treatment of points made in his earlier work (36). Fuchs (45) argues along a somewhat similar line that the correct notion of killing in the context of physician action is "teleological" or "biological," and that this notion resists being reduced to equivalence with letting die.

[36] Huddle concurs: "It is likely, perhaps, that fault lines in this debate go deeper than any conceptual analysis of doing and allowing can bridge. The kinds of moral theory friendly to the doing-allowing distinction or to its elimination track, more or less, deontology and consequentialism. The gulf between these approaches to moral theory remains wide and we should not, perhaps, expect a resolution of debate over the doing-allowing distinction sooner than that broader debate is resolved – if, indeed, we should expect any such resolution" (55, p.262).

Attempting to answer question 3 by first providing a (general) answer to 2 or (*a fortiori*) to 1, and then applying that general answer to 3 (treating 3 as an instance of 2 or 1) is an exercise in abstraction. Depending on how the general answer to 1 or 2 is formed (how it is derived from reflection on particular cases or from conceptual analysis), its application to 3 may or may not be helpful. A generalization's application to the clinical situation will be unhelpful when the particulars from which the generalization was derived either contain aspects that differ importantly from the clinical situation or lack important aspects of the clinical situation. There are many forms of doing and killing that are unlike the doing/killing involved in VAE and, more importantly, there are many forms of allowing, and allowing to die, that are unlike the allowing to die involved in WLT. Generalizations drawn inductively from cases that differ in significant ways from the clinical situation may, for that reason, suffer from irrelevance to the comparative assessment of VAE and WLT. Similarly, generalizations formed by analysis of concepts (e.g., "doing" and "allowing"; "killing" and "allowing to die") may fail to apply helpfully to the assessment of VAE and WLT if the conceptual analyses take place without regard to key features of the clinical situation.

Negative Examples: The Tendency to Irrelevant Abstractions

James Rachels famously pioneered the "contrast strategy," discussed briefly in Part A as one of the objections to the Doing/Allowing Account (7). The contrast strategy purports to show that the distinction between killing and allowing to die is irrelevant to morality, by means of a comparison of cases. Accordingly, if we can construct two cases that are identical in every way except that one is a killing and the other an allowing to die, and yet we judge that the one who allows to die is just as morally culpable as the one who kills, then killing (in itself) does not differ morally from allowing to die (in itself) (5-10,56,57).[37] Woollard (16, p.461) follows Kagan (58) in arguing that the contrast strategy is illegitimate (whether used to attack or defend the moral significance of the distinction) because it "wrongly assumes that if a factor makes a moral difference anywhere, it will make the same moral difference everywhere".[38] While Rachels' (7) cases support the intuition that killing and allowing to die are morally equivalent, comparison of other pairs of cases can elicit the opposite intuition, that killing and allowing to die are *not* morally equivalent (16,55,58).[39] Generalizing from Rachels' cases, we would conclude that the distinction is not morally significant, but since generalizing from other paired cases yields the opposite conclusion, the contrast strategy cannot establish a universal proposition in either direction.

Given the problem with the contrast strategy just observed, we should ask whether there are some cases that serve better than others to form the basis for generalizations that are applicable to the clinical situation (the situation that includes both VAE and WLT as clinical practices). Are there specific features of the clinical situation that need to be reflected in any such general proposition about killing and allowing to die for that proposition to apply helpfully to a comparative assessment of VAE and WLT? When the expression "allowing to die" is used in the context of a health care professional's withdrawing of life support at the request of a competent patient (call this "context W"), it carries significant connotations and nuances that it does not have in the relevant Rachels cases.

From our review of the literature in part one of this paper, and especially by attending to features noted in accounts that are more sensitive to clinical contexts, we have derived the following list of distinctive, ethically relevant aspects of the clinical situation. In context of use W, "allowing to die" standardly connotes that,

a) the act of WLT is taking place within an established and ongoing fiduciary relationship between a health care professional and a patient;
b) the health care professional has taken responsibility for the patient's care by providing life-sustaining treatment;
c) there is an underlying illness that would have resulted in the patient's death, were it not for the life-sustaining treatment;
d) the health care professional's activities are bounded by the patient's negative right to be free from unwanted bodily invasion;
e) it is possible to engage in the act without intending the death of the patient;
f) the patient, while competent and not under duress or pressure, has voluntarily requested that treatment be stopped.

Relatedly, many similar aspects would hold in the context of a health care professional performing VAE by administering a lethal injection at the request of a competent patient (call this "context E"). Either as worded or in a slightly modified form,

[37] In Rachels' paired cases Jones, intent on securing an inheritance, allows his young nephew to drown in a bathtub, while Smith, with the same intention, kills his nephew by drowning him (holding his head under the water). Smith acts by commission (kills), while Jones acts by omission (allows to die). But, according to Rachels, these facts are irrelevant to the moral assessment of the two cases because both are equally guilty of murdering their nephew. Thus, the distinction between killing and allowing to die is not morally significant. Tooley and Lichtenberg use different case comparisons to argue in a similar way (8,9). Brock and Quill (5) and Lo and Brock (6) use the contrast strategy to argue for the moral equivalence of a doctor who withdraws treatment and a greedy nephew who removes the same treatment. A similar use of the strategy occurs in Orentlicher (56) and Billings and Block (57) concerning cases of terminal sedation in comparison with VAE.

[38] Huddle (55) provides a related critique that specifically takes aim at the use of comparison cases in Causation Accounts.

[39] Kagan (58) considers the cases of allowing someone trapped on a hillside to be crushed by a rolling boulder in order to reach the hospital in time to save the lives of two companions versus pushing the boulder so that it crushes someone trapped on a hillside in order to open the route to the hospital to save the lives of the companions. Our intuitions tend to incline us to see these cases, which differ only in that one involves killing and the other allowing to die, as morally different.

aspects (a), (b), (c), (d) and (f) also obtain in context E.[40] Arguably, these connotations or nuances of meaning and context are not factors extraneous to the issue of the killing / letting die distinction; they are partially constitutive of the meaning of those expressions *in the context that is relevant to the debate about WLT and VAE* (59).[41] Note that we do not claim that (a) through (f) exhaust the ethically relevant aspects unique to the clinical situation that might encompass contexts W and E.

What it means for a health care professional to "allow (their patient) to die" differs from what it means for Smith to "allow (his nephew) to die," because the act specified by the expression "allow to die" in the different contexts is carried out by different agents who stand in different relationships to the patients (those on whom the act is carried out), and hence in different nexuses of responsibility (30,31,34). Given Jones' relationship with his nephew, Jones has a duty of care that mandates that he act to rescue his nephew. By contrast, when a competent patient requests WLT, this removes the physician's responsibility to continue to provide the treatment. Jones' intention (regardless of whether intentions are constitutive parts of actions) makes Jones' act immoral – Jones intends the child's death as a means to secure an inheritance. By contrast, a physician need not intend their patient's death. So even though the physician, like Jones, is in an established and ongoing relationship of care toward another (connotation a) and has taken responsibility for that other's care (connotation b), and there is an underlying condition that will lead toward death unless intervention occurs (connotation c), connotations d and e do not hold in Jones's case. As an instance of allowing to die, then, Jones's case fails to sustain the sort of generalization that would meaningfully apply to standard cases of WLT.

Isaacs (51) illustrates another way conceptual analysis can go astray when carried on in abstraction from the concrete clinical situation. Isaacs carries out analysis at the level of question 2 above: Is there a moral difference between causing someone's death (making them die) and merely allowing them to die? The purportedly significant difference Isaacs considers is that "killing is, in some sense, worse than letting die" (51, p.357). The argument is that since the distinction between killing and letting die is either aligned with intuitive moral judgments but unclear, or clear but out of touch with moral judgments, the distinction is not useful for moral deliberation and is therefore not morally significant. But Isaacs supports the claim that clear accounts of the distinction fail to get the right intuitive results with only a single illustration of a clear account: *If* killing is an action and letting die an inaction, and *if* actions are bodily movements, then both WLT and VAE are killings (contrary to our intuitions).[42] More to the point, Isaacs addresses the other horn of the dilemma – that when drawn in accordance with common judgments about cases the distinction is unclear or illegitimate – by considering McMahan's (31) complex system of criteria. McMahan's account of the distinction is based on cases like the Dutch Boy who gets tired and pulls his finger from the dike (allowing the townspeople to die) and his father who yanks his son's finger out (killing the townspeople), and the Firefighter who moves his safety net away from one falling individual to rescue two others. Isaacs complains that the distinction has now become unwieldy and confusing because of all the complexity packed into how it is characterized. We suggest that the problem is rather that too much is being expected of the distinction, drawn at that level of generality.

In order to cover so many disparate sorts of cases, the conceptions of killing and letting die have to be conceptually thinned out, abstracting away from important aspects of concrete situations like the clinical situation in which WLT and VAE occur. To take one instance, the relationship between the Dutch Boy and the townspeople he is taken to have let die differs significantly from the relationship between a doctor and their patient. Connotations a, b and d from Context W do not hold in Dutch Boy, and though connotation c does hold, it holds in a meaningfully different way.[43] But because of these differences, the Dutch Boy's "letting die," although it falls under the general concept of removing an obstacle that was preventing an outcome, is a different sort of act from WLT. So, contra Isaacs, McMahan's adaptation of the distinction between killing and letting die does not rob the distinction between WLT and VAE of its moral significance. Instead, the problem is with the assumption, made by both McMahan and Isaacs, that such a general conception of killing and letting die captures the sense of those terms in the clinical context. Since, as shown above, it *does not* capture that sense (because it abstracts away from key aspects of the clinical context), claims about the generalized distinction's clarity or unclarity, or its alignment or non-alignment with moral judgments, are neither here nor there concerning the comparative assessment of WLT and VAE.

[40] In situations where VAE takes place in the context of an established patient to health care professional relationship, and substituting "VAE" for "WLT," aspects (a) and (f) would obtain straightforwardly in context E, and (b) would hold in a modified form since the professional in question will have taken responsibility for patient care but that care may or may not involve provision of life-sustaining treatment. In jurisdictions where having a "grievous, irremediable illness" is a condition for legitimate practice of VAE, a form of (c) would hold in context E. Some will want to contest the ethical relevance of aspects (d) and (e) since that relevance depends in part on views of the importance of certain rights (aspect d) and intentions in relation to the practice of WLT and VAE (aspect e).

[41] Arguing against the contrast strategy championed by Rachels and others (7-10), Asscher (59, p.272) takes aim at the use of the "Similarity Criterion," which holds that "cases must be exactly alike, except for the moral issue at hand and excepting differences that clearly have no moral relevance." The Similarity Criterion is used to produce examples that avoid extrinsic differences (such as differences in responsibility, motivation, or intention), as cases without extrinsic differences, such as Rachels' *Bathtub*, purport to give special insight into the distinction itself (i.e., they show what, if any, *intrinsic* differences exist between killing and allowing to die). Asscher views the Similarity Criterion as a faulty philosophical tool because the examples are problematic in themselves (they always contain some extrinsic difference and, if avoided, create problematic similarities or gaps in the cases that skew moral conclusions) and its use precludes counterexamples merely because they contain extrinsic differences. Asscher holds that debate regarding the distinction has not progressed due to treating such similar cases as paradigmatic, but by rejecting the Similarity Criterion and its cases, alternative accounts (containing explanatory extrinsic differences) are allowed to develop, thereby enabling the debate to advance. We consider Asscher's own positive account in the following section.

[42] The second "if" in the previous sentence is perhaps bigger than the first; but since only one possible theory of action, in combination with one way of drawing the distinction have been considered, Isaacs can hardly conclude that there is no clear account of the distinction that accords with our moral judgments.

[43] The Dutch Boy is not in an established relationship with the townspeople (not-a), nor has he taken professionally or contractually binding responsibility for their protection from natural disasters (not-b), nor are the townspeople requesting to be allowed to die so that their rights to be free from unwanted bodily invasion would be violated if the boy refused their request (not-d). There is a threat (the compromised dike) that, were it not for the boy's finger-plugging, would result in the death of the townspeople, but this threat is not one for which the boy can be held (professionally or contractually) responsible for averting (so not-really-c either).

Finally, Atkinson's (60) conceptual analysis is more satisfying than many other case-based analyses (7-10), while still suffering from problematic abstraction when applied to VAE and WLT as clinical practices. The careful descriptive work and examination of several diverse cases of killing and allowing to die makes Atkinson's approach a satisfying way of crafting a distinction between them at a general level. But as a result the distinction loses its grip on the clinical context.[44] The cases sampled range from agents providing a lethal injection or terminal sedation, withholding or withdrawing life-sustaining treatment, and engineering or bringing about death by opportunistic means.[45] To be sure, not all the cases or their treatment involve the key features of the clinical setting relating to WLT and VAE, and to that extent a generalization drawn from them will be problematically abstract in relation to an assessment of those practices.

To better understand where the problem of abstraction occurs for Atkinson, consider the conclusions drawn from his examination of cases. Atkinson contends that if the distinction is used as a *moral distinction*, then the terms of the distinction already contain a prior moral evaluation – namely, killing is morally worse than allowing to die. The normative components in the distinction might be morally relevant, but to draw any conclusions by using the distinction so characterized begs the question. If the distinction is used as a *descriptive distinction*, then nothing of moral significance follows from the distinction itself, as there are other moral features of a case that can do the operative work – the distinction itself is not sufficient to ground any moral difference. Atkinson discusses a difference in intent, long-term social consequences, cause of death, the reversibility of error, responsibility, and a difference in degree of interference as some features that are independent of the distinction itself that can do the operative moral work.[46] He suggests that what makes a moral difference in one case may not hold in another, such that critical analysis is always required. As such, it follows from Atkinson's analysis that the distinction on its own is not useful for general moral deliberation.

It is important to reiterate that Atkinson's goal was not to characterize the distinction as applied to a clinical setting, but instead to conceptualize and evaluate the distinction for moral reasoning *in general*. However, Atkinson's approach interestingly shows a way to see how examining WLT and VAE *concretely* (i.e., examining the features of the clinical setting) may reveal a moral difference that *is* useful for practical moral deliberation in its context of use. An account of the distinction that focuses on the clinical situation and its distinctive features may not hold in other contexts that involve cases of killing and allowing to die, and so there may not be an account that covers all situations. But while Atkinson is surely right about this, to conclude that no adequately specified forms of the distinction between killing and allowing to die can be useful would be a mistake. The distinction, as characterized by Atkinson's general analysis, is largely irrelevant to a comparative assessment of VAE and WLT, but this simply points to the need to formulate more carefully what killing and allowing to die involve when they take place within concrete contexts like the clinical setting and relational-obligational matrices in which VAE and WLT take place.

To sum up, as illustrated in this section, discussions of the moral significance of the distinction between WLT and VAE that treat these clinical practices as instances of killing and letting die (conceived generally) tend to be unhelpful for the purpose of making an adequate comparative assessment of the practices. Similarly, argumentative strategies that fail to reflect the way distinctive aspects of the clinical situation contribute to a determination of the acts being compared tend toward irrelevance, as far as an adequate comparative assessment of those acts is concerned.

Positive Examples (Concrete Treatments)

The previous section described some representative examples of approaches that veer toward irrelevance due to the inapplicability of generalizations about killing and allowing to die to the consideration of VAE and WLT in the concrete clinical situation. But many authors engaged in the debate proceed in ways that are more sensitive to key features of that situation. We turn now to three examples of more sensitive, and therefore more satisfying approaches. Asscher's (34) Responsibility Account is compelling in large part because what it says about the distinction between killing and letting die turns on the responsibility that accrues specifically to healthcare professionals, which in turn is affected by three key factors: 1) the nature of the relationship they have entered into with their patients, 2) the choices their patients have expressed, and 3) the surrounding responsibility of the institution they work for.[47] On Asscher's account, if a patient requests active treatment and a healthcare team accept responsibility for treatment, to let the patient die (by withdrawing treatment) would be as morally wrong as to kill the patient (by administering a lethal injection). In that case, killing and letting die both violate the responsibility to provide active treatment. If the patient instead asked that active treatment be discontinued and that only palliative care be given, then it would clearly be morally worse to kill the patient than to allow them to die. In such a case, letting the patient die would be permissible since there is no responsibility to provide active treatment. A patient whose life is worth living who approaches a healthcare team asking to be killed would be harmed if killed by a member of that team, and for that reason it would be worse to kill them than to allow them to die (by refusing the patient's request and letting them walk

[44] Atkinson's general characterization of the distinction states: 1) Instances of killing involve a causal process where an agent initiates, redirects, restructures, or supports a process leading to death, 2) Instances of allowing to die involve some independence of an agent to the causal process, and 3) Killing is typically worse than allowing to die, such that "the term 'killing' carries a more strongly negative tone than does 'letting die'" (60, p.1922). With respect to context W and as applied to WLT, this characterization of allowing to die only satisfies connotation (c).

[45] For example, not turning on the heat in a person's room such that they catch pneumonia and die or, in a structurally similar case to Rachels' *Bathtub*, "Planning to turn off the heat and so cause his death, I drop in on Jones only to find that the heat has already accidentally been turned off and so experience relief in having the matter taken out of my hands" (60, p.1915). From our discussion of Rachels' argument above, comparing agents who stand in different relationships and nexuses of responsibility to the patients may not be helpful to a comparative moral assessment of VAE and WLT as clinical practices.

[46] From the list of features Atkinson describes, some features directly relate to context W, namely, responsibility (connotation b), cause of death (connotation c), and difference in intention (connotation e).

[47] 1) correlates with connotation (a) from Context W, 2) correlates with connotation (d) and 3) corresponds to connotation (b).

out). But on the assumption that a patient's life is not worth living, Asscher asserts, assuming that the healthcare team had taken responsibility for the patient's care, it could be better, morally speaking, to kill the patient who requests to be killed than to allow them to die.[48] Whatever we think about Asscher's success in accounting for a morally significant difference between WLT and VAE, we can appreciate that his approach is helpful for the purpose of assessing these practices. What makes it helpful is that it explicitly considers and compares the practices in light of distinctive features of the clinical situation rather than under the assumption that such features are unimportant and can be ignored.

Paterson's Cluster Account (27), in refusing to deal with factors in isolation, affords another satisfying account of the difference between killing and allowing to die in the clinical setting. For Paterson, the judgment that a practice involves wrongful killing (homicide) depends on the interworking of multiple factors. When the patient has an underlying terminal pathology, *and* a doctor (or other healthcare professional) acts with the intention to bring about the patient's death, *and* this act violates an existing moral duty, then we should judge this as a case of wrongful killing. The doctor could so act either by commission (by giving a lethal injection) or by omission (by removing life-support); depending on the moral duties accruing to the doctor in the situation, the doctor might kill the patient by withdrawing life-support just as much as by administering an injection. But for Paterson, an action may constitute an act of homicide even without the doctor explicitly intending the patient's death; the doctor might not have intended the patient's death, but if the patient's death is a side-effect that has not been justified through an adequate assessment of possible results that a responsible physician should have carried out, the act may still involve wrongful killing. Only when the doctor has not acted with the intent to kill *and* has not culpably avoided engaging in an appropriate justification of the patient's death as a side-effect can the doctor's act be judged as an instance of "letting die". In such cases, Paterson reasons, the doctor is not morally culpable, the proximate cause should be attributed to the underlying pathology and the doctor's action treated as a secondary, non-determinate cause of the patient's death (27). Again, our goal is not to promote Paterson's account of the distinction between killing and letting die but to point out that in some important ways Paterson's approach to the issue improves on the approaches taken in the previous section.[49] In particular, this approach attempts to reflect carefully the relevance of diverse and distinctive aspects of the clinical situation (fiduciary doctor-patient relationship (a); responsibility for treatment (b); underlying critical illness (c); patient rights (d); clinician intention (e); voluntariness of request (f)) to an adequate comparative assessment of the practices of WLT and VAE.

Finally, the account collectively provided by Miller, Truog, and Brock (22,28) differs from the other examples in the present section in not arguing directly for or against the moral relevance of the distinction between killing and allowing to die but contending instead that WLT is a species of 'killing' rather than 'letting die.' They oppose using the term 'killing' as it is commonly understood to mean "the unjustified taking of life" in medical contexts, preferring instead to define 'killing' as 'causing death' since such causing "may or may not be morally justified, depending on the circumstances" (22, p.457). By using this approach, they aim to discuss medical conduct associated with death without the negative moral connotations associated with the term 'killing' (28). It follows from this definitional point that if WLT is accepted to cause death, then the traditional medical ethics perspective of treating WLT as morally distinct from VAE is mistaken.

What makes Miller, Truog, and Brock's account satisfying is that their approach represents an attempt to examine the metaphysics of action in the clinical situation of cases involving WLT. They contend that an honest examination of WLT in context shows that this routine medical conduct causes death and clearly conflicts with the traditional moral norm that doctors may never kill (22,28). Using a pair of contrasting cases, they illustrate inconsistencies with the moral judgments between WLT and VAE by examining the conventional normative responses towards suicide, causation, intention, and responsibility ascriptions (22). Contrary to more traditional understandings,[50] they claim that:

1. WLT and VAE are cases requesting suicide, as suicide can be understood as aiming at causing one's own death;[51]
2. WLT contributes to causing death in a relevant way, as the withdrawal makes a difference in explaining the timing of death in some ethically legitimate cases;[52]
3. If it is accepted that WLT causes death and is ethically legitimate, then the ethical legitimacy of the conduct remains whether a physician intends to cause death or not;[53]

[48] It follows that authors who assume that the moral difference between killing and letting die consists in killing being morally worse than letting die have oversimplified the situation. Killing and letting die, in the context of a comparison of VAE and WLT, can differ in ways relevant to moral assessment that go beyond the mere judgment that killing is always wrong (morally unjustified) while letting die is always right (morally justified). Similar points were made in footnotes 7, 27 and 28 above.

[49] For instance, for the purposes of this paper, we have remained neutral concerning the correct sense of the term "intention" and its place in accounts of killing and allowing to die, while Paterson's account relies heavily on a certain understanding and prioritization of the term.

[50] Recall that the standard view in respect to the Causation Account maintains that WLT does not cause death, but merely allows a patient to die from their underlying medical condition. Furthermore, a traditionally and strongly held moral norm asserts that it is impermissible for doctors to *intentionally* cause death.

[51] Relatedly, physician assistance in these requests is thus assisted suicide.

[52] For example, those who are "not terminally ill and [have] the potential to live at least 10 years while being maintained on a ventilator," are cases where if it were not for WLT, then the patient would still be alive (22, p.453). As noted earlier, the account of Miller et al. depends on understanding causation as difference-making, which may or may not be the right notion of causation for this context.

[53] They argue that the claim that physicians *never do* intend to cause death is false, and furthermore, that claiming physicians *never should* intend to cause death is also suspect, as a physician can intend to respect a patient's choice for WLT and accept that they cause death (intentionally) in doing so. They do not claim the reverse: that clinicians *always* intend to cause death with WLT. As noted earlier, this point is questionable depending on how the term "intention" is operationalized.

4. Causation does not equate to culpability, as responsibility for causing death can be attributed without holding physicians culpable for wrong doing (i.e., causing death by WLT, or other means, can be right or wrong depending on the circumstances, such as in cases of negligence or whether informed consent is provided by competent patient or not).

To diagnose the inconsistency between their assessment and the standard medical ethics view, they propose the concept of *moral biases* (28).[54] A moral bias is "a *motivated* false belief about human conduct that serves a *legitimating* function" (28, p.9). Miller et al. recognize that there is a strong motivation in clinical settings to interpret practices causing death as something other than killing; the legitimating function is to remove the conflict between routine conduct (that causes death) and the prevailing moral norm that prohibits such conduct.[55] In other words, moral biases essentially distort the metaphysical account of the conduct, or how things *are*, to cohere with the established norms dictating how things *should* or *should not* be.[56] More simply, moral biases alter facts about WLT to eliminate conflict with the norm that doctors must not cause death. Miller, Truog, and Brock's account attempts to explain why clinicians and observers tend to *feel* like WLT causes death, namely because it *does* cause death, despite the conventional view denying the causal connection between WLT and death (28,61). They conclude that the routine practices can still be justified when moral biases are abandoned and the ethically suspect norm is reconstructed, as other traditional ethical principles and norms can continue to justify them (e.g., non-maleficence & beneficence). However, it follows that one must then give up the differential moral assessments between WLT and VAE.

With respect to the key connotations of context W, Miller, Truog, and Brock's account is an attempt to reconstruct how assessments of VAE and WLT are conducted. Consider their contrasting case of John and Sam, two persons living with similar spinal cord injuries from similar accidents. John is ventilator dependent while Sam now breathes spontaneously and no longer requires a ventilator, both do not want to continue living and cite the same reasons for wanting to die peacefully.[57] One requests that their physician withdraw the ventilator (WLT) while the other requests their physician to administer a lethal dose of medication (VAE). Here, connotation (a) holds in both cases and, it may be argued, some version of connotation (b) also holds (62).[58] Due to WLT contributing to causing death in a relevant way, connotation (c) is mistaken, as the patient would still be alive were it not for WLT; this shows the way in which connotation (c) may be irrelevant in other cases. Connotation (d) holds for John and not for Sam. The right to be free from life-sustaining treatment is no longer relevant, but there is another right to consider for the purposes of comparing WLT and VAE – namely, the right to non-interference with the action of a willing physician. While responsibility for recognizing this right is beyond individual physicians, the suggestion is that connotation (d) holds for WLT but is irrelevant in comparing WLT and VAE. Finally, connotation (e) standardly holds for John and not for Sam. However, Miller et al.'s analysis (in general) calls connotation (e) into question. They contend that the relevance of intention, as evoked in connotation (e), only obtains if death is viewed as an absolute harm in all circumstances; if death is not viewed in this way, then connotation (e), which concerns physicians' intentions, may also be irrelevant in comparing WLT and VAE. Since intentions only matter when conjoined with the absolute prohibition against causing patient death, (e) will be irrelevant if WLT and VAE are both ethically acceptable exceptions to that prohibition (as Miller et al. take them to be).[59] In summary, Miller et al.'s account, which contests divergent moral assessments of WLT (context W) and VAE (context E), indicates possible disagreement over the ethical relevance and application of the key clinical features and signals how fruitful debate in this area can be accomplished when paying attention to these features.

In this section, we have shown how approaches that explicitly treat key features of the concrete clinical situation contribute more effectively to the comparative assessment of WLT and VAE. Despite the fact that the three approaches considered here differ in what claims they make about the relation of WLT and VAE to the concepts of "killing" and "allowing to die," they all promote fruitful debate in this area by keeping an attentive eye on relevant features of the clinical situation. In attempting to answer the question of whether there is a morally relevant difference between WLT and VAE, these are the sorts of approaches upon which we need to focus, just as we should avoid the sort of approaches shown in the previous section to suffer from irrelevant abstraction. Accounts that focus on the clinical situation and its distinctive features may not help with comparisons of other forms of killing and allowing to die (the difference between, say, not sending aid to save the life of a starving child and killing a starving child who is trying to take money from your wallet). But if our assessment holds, *those* distinctions will be best addressed by focusing on examples that preserve all the distinctive features of *those* situations.

[54] Earlier, they refer to these as "moral fictions" (22, p.454).
[55] They also describe moral biases as tools to cope with the "cognitive dissonance" faced from the conflict between the accepted practice and the apparently conflicting moral norm (22,28). Isaacs makes a very similar point (51).
[56] In referring to how moral biases operate, Miller and Truog (28, p.20) state that they "may be seen…as a reversal of the so-called 'naturalistic fallacy,'" that is, moral predicates influence non-moral (or metaphysical) predicates.
[57] It is assumed that John, like Sam, finds his life independently burdensome such that his interest in ending his life does "not stem from the mere fact that he needs assistance of a burdensome mechanical ventilator to breath" (22, p.455).
[58] For example, in this case, the fact that life-sustaining treatment is no longer in place for Sam does not remove the health care professional's responsibility taken on when life-sustaining treatment was initiated in the first place. The responsibility for the surrounding situation (i.e., a life independently burdensome as a result of the initiation of life-sustaining treatment) holds for John and Sam, such that it would be unjust to treat Sam differently merely because he no longer requires a ventilator. Beauchamp neatly summarizes this point: "Medicine and law seem now to say to many patients, 'If you were on life-sustaining treatment, you could withdraw the treatment and we could let you die. But since you are not, we can only give you palliative care until you die a natural death.' This position condemns the patient to live out a life he or she does not want – a form of cruelty that violates the patient's rights and prevents discharge of the fiduciary obligations of the physician" (62, p.82).
[59] Of course, whether VAE and WLT should be seen as ethically acceptable exceptions to the prohibition against causing patient death is a centrally important and contested point. Furthermore, recall the disagreement between Sulmasy and Miller et al. in our discussion of the Intention account. On Sulmasy's account, connotation (e) is relevant to WLT because intentions can vary in carrying out WLT. A physician may intend (but may or may not be committed to) the patient's death in WLT, whereas VAE always involves the intention to kill the patient. So, connotation (e) obtains if intentions are understood along the lines of Sulmasy's account. But if Sulmasy is wrong to understand intentions that way, then connotation (e) won't differentiate between WLT and VAE.

RECOMMENDATIONS FOR MOVING THE DEBATE FORWARD

Features of the general concepts "killing" and "allowing to die" do not establish or refute a morally relevant distinction between withdrawing life-sustaining treatment and administering a lethal injection. Intuitions and generalizations derived from paired cases involving situations that are significantly removed from real clinical practice tend to be unhelpful in addressing the euthanasia debate. General conclusions about killing and letting die reached on the grounds of examples that ignore key aspects of the clinical situation tend to be irrelevant to a comparison of the clinical practices (WLT and VAE) we set out to examine. Arguably, doctors who have claimed that ending a patient's life by a lethal injection is morally distinct from removing a patient's life support never meant to locate the moral distinctness in an abstract examination of concepts. To discover whether the distinction is morally significant, we must attend to the difference between the practices *in concreto*. For this reason, treatments of the distinction as it applies to the specific practices tend to be much more satisfying and adequate, as shown above. (At the very least they get us into the right ballpark, so to speak).

A more empirically- and contextually-informed approach to the debate is needed. To discern the moral significance (or lack thereof) of the distinction, we need to "get inside the heads" of the moral agents actually involved in the practices under consideration. Do doctors find accounts of the distinction between killing and allowing to die (in terms of doing vs. allowing, causal responsibility, intentions, etc.) helpful in thinking through whether they experience the practices as different? The approach we wish to propose would not directly determine whether there is or is not an actual difference, from a normative standpoint. Instead, it would consider that if the distinction is morally relevant (makes a moral difference), then it must be useful for moral deliberation by real moral agents. We assume that a distinction's relevance to (or usefulness for) moral deliberation and action counts in favour of its moral significance, since such relevance (or usefulness) is at least part of what it means for a distinction to be morally significant (51).[60] To put this point more strongly, we believe that such relevance constitutes a necessary condition for moral significance. So, if physicians experience WLT as equivalent to VAE despite an awareness of the way the distinction between the two practices can be drawn, that would suggest that the bioethical distinction *does not* make a difference for practical action in the clinical setting. And if the distinction can be shown to be irrelevant or useless for the moral deliberation of real moral agents, that would strongly suggest that the distinction lacked moral significance.[61]

To advance this debate, clinicians' experiences, and whether they perceive or fail to perceive (moral) differences between the practices, should be taken seriously. Attending to these experiences and perceptions will help us to formulate more nuanced positions on the distinction between killing and allowing to die. Exploring physician experience can reveal how the distinction (as a tool) may or may not be useful, and provide the needed context to theorize in a more contextually informed, and therefore more practically useful, way.

We also recommend listening for additional features of the clinical situation that doctors would identify as relevant to their own understandings and assessments of the practices. Again, we believe such exploratory research to be warranted, if not required, by our discussion above of the way approaches to the debate can be more or less helpful. Progress in this debate will be facilitated by sensitivity to features of the clinical situation that must be reflected for an adequate understanding of the acts being compared. Without such an understanding, we cannot expect to get very far with the comparative moral assessment of the practices.

CONCLUSION

To be clear, we do not anticipate that an empirical approach of the sort we propose here will, on its own, resolve the ethical debate. An empirical study of physicians' experiences will not tell us what is in fact right and wrong, from a normative standpoint. But since, as we have observed, the more satisfying approaches to the question are those that attend carefully to distinctive features of the clinical situation rather than assuming that an abstract analysis will transfer over to the relevant clinical cases, it is reasonable to think that attending to the experience of physicians will help identify what features of the situation must be considered in any adequate ethical analysis. Furthermore, attending to physicians' experience of the practices can help us to determine whether the distinction matters for practical action in the clinical setting. To get closer to the heart of the issue, we need to listen to doctors who engage in both VAE and WLT, and to their own accounts of the moral difference, or lack thereof, between the practices we want to compare.[62]

[60] At the very least, a distinction that systematically failed to aid the moral deliberation of agents in the very situations to which that distinction purportedly applied would, *ceteris paribus*, be reasonably regarded as lacking moral significance.

[61] Logically, if usefulness for practical moral deliberation is a necessary condition for moral significance, then by *modus tollens*, to show that the distinction is not useful for moral deliberation would demonstrate that the distinction is not morally significant. Conversely, to show that the distinction is useful for moral deliberation would not be to show that it is morally significant (that would be to affirm the consequent), but it would establish a necessary condition for moral significance, and would thereby support, to a greater or lesser extent (depending on what else is needed to make up a sufficient condition for significance), the claim that the distinction is morally significant. To be fair, we must note that whether doctors say, or even sincerely believe, that they experience the two practices as distinct or not does not necessarily establish that the distinction is or is not useful. After all, the prior convictions and values of the doctors concerning life, death, suffering, compassion, etc. may influence the way they interpret their own experiences in one way or the other. What doctors say about their experience might involve a (conscious or unconscious) attempt to justify their fundamental convictions and values, rather than being an objective report of that experience. It will require some substantial interviewing and interpretive skill to be able to say reliably, on the basis of the doctors' reports, whether the distinction is or is not useful for these agents' moral deliberation. Thanks to Lauris Kaldjian for raising this important issue for the approach we are recommending.

[62] We intend to conduct interviews with physicians in 2020 and to interpret and publish the results in 2021, as part of a research project funded by the Insight Development Grant stream of the Social Sciences and Humanities Research Council of Canada.

Reçu/Received: 10/07/2020
Publié/Published: 1/12/2021

Remerciements
Cette recherche a été soutenue par une subvention de développement Savoir du Conseil de recherches en sciences humaines du Canada et par des fonds de recherche équivalents fournis par l'Université St-Jérôme. Nous tenons à remercier les personnes suivantes pour leurs commentaires utiles sur les versions préliminaires de ce document : Kieran Bonner, Lydia Dugdale, Ewan Goligher, Lauris Kaldjian, Daniel Kim, Jane Kuepfer, Midori Matthew, John Neary, Nick Ray, Sylvia Terzian, Donna Ward, John Yoon.

Acknowledgements
This research was supported by an Insight Development Grant from the Social Sciences and Humanities Research Council of Canada, and by matching research funds provided by St. Jerome's University. We would like to thank the following individuals for their helpful feedback on drafts of the paper: Kieran Bonner, Lydia Dugdale, Ewan Goligher, Lauris Kaldjian, Daniel Kim, Jane Kuepfer, Midori Matthew, John Neary, Nick Ray, Sylvia Terzian, Donna Ward, John Yoon.

Conflits d'intérêts
Aucun à déclarer

Conflicts of Interest
None to declare

Édition/Editors: Nico Nortjé & Aliya Affdal
Les éditeurs suivent les recommandations et les procédures décrites dans le Code of Conduct and Best Practice Guidelines for Journal Editors de COPE. Plus précisément, ils travaillent pour s'assurer des plus hautes normes éthiques de la publication, y compris l'identification et la gestion des conflits d'intérêts (pour les éditeurs et pour les auteurs), la juste évaluation des manuscrits et la publication de manuscrits qui répondent aux normes d'excellence de la revue.

The editors follow the recommendations and procedures outlined in the COPE Code of Conduct and Best Practice Guidelines for Journal Editors. Specifically, the editors will work to ensure the highest ethical standards of publication, including: the identification and management of conflicts of interest (for editors and for authors), the fair evaluation of manuscripts, and the publication of manuscripts that meet the journal's standards of excellence.

Évaluation/Peer-Review: Johan Bester & Barbara Pesut
Les recommandations des évaluateurs externes sont prises en considération de façon sérieuse par les éditeurs et les auteurs dans la préparation des manuscrits pour publication. Toutefois, être nommé comme évaluateurs n'indique pas nécessairement l'approbation de ce manuscrit. Les éditeurs de la *Revue canadienne de bioéthique* assument la responsabilité entière de l'acceptation finale et de la publication d'un article.

Reviewer evaluations are given serious consideration by the editors and authors in the preparation of manuscripts for publication. Nonetheless, being named as a reviewer does not necessarily denote approval of a manuscript; the editors of *Canadian Journal of Bioethics* take full responsibility for final acceptance and publication of an article.

REFERENCES

1. Chandler J. Killing and letting die—Putting the debate in context. Australasian Journal of Philosophy. 1990;68(4):420-431.
2. McGrath S. Causation and the making/allowing distinction. Philosophical Studies. 2003;114(1-2):81-106.
3. Foot P. The problem of abortion and the doctrine of double effect (1967). In: Steinbock B, Norcross A, eds. Killing and Letting Die. 2nd ed. New York: Fordham University Press; 1994. p.266-279.
4. Bronner B. Two ways to kill a patient. Journal of Medicine and Philosophy. 2018;43(1):44-63.
5. Brock DW. Voluntary active euthanasia. In: Beauchamp TL, Walters L, eds. Contemporary Issues in Bioethics. 5th ed. California: Wadsworth Publishing; 1999. p.296-305.
6. Quill TE, Lo B, Brock DW. Palliative options of last resort: comparison of voluntarily stopping eating and drinking, terminal sedation, physician assisted suicide, and voluntary active euthanasia. JAMA. 1997;278(23): 2099-2104.
7. Rachels J. Active and passive euthanasia. The New England Journal of Medicine. 1975;292(2):78-80.
8. Lichtenberg J. The moral equivalence of action and omission. Canadian Journal of Philosophy. 1982;8:19-36.
9. Tooley M. An Irrelevant consideration: killing versus letting die. In: Steinbock B, Norcross A, eds. Killing and Letting Die. 2nd ed. New York: Fordham University Press; 1994. p.103-111.
10. Singer P. Practical Ethics. 2nd ed. Cambridge: Cambridge University Press; 1993. p.206-213.
11. Begley AM. Acts, omissions, intentions and motives: A philosophical examination of the moral distinction between killing and letting die. Journal of Advanced Nursing. 1998;28(4):865-873.
12. von Wright GH. Acta Philosophica Fennica: An Essay in Deontic Logic and the General Theory of Action. Amsterdam: North Holland Publishing Co; 1968.
13. Walton DN. On Defining Death: An Analytical Study of the Concept of Death in Philosophy and Medical Ethics. Montreal: McGill-Queen's University Press; 1979.
14. Merkel R. Killing or letting die? Proposal of a (somewhat) new answer to a perennial question. Journal of Medical Ethics. 2016;42(6):353-360.
15. McLachlan H. The ethics of killing and letting die: active and passive euthanasia. Journal of Medical Ethics. 2008;34(8):636-638.
16. Woollard F. The doctrine of doing and allowing II: The moral relevance of the Doing/Allowing distinction. Philosophy Compass. 2012;7(7):459-469.
17. Foot P. Killing and letting die. In: Steinbock B, Norcross A, eds. Killing and Letting Die. 2nd ed. New York: Fordham University Press; 1994. p.280-289.

18. Quinn WS. Actions, intentions, and consequences: the doctrine of doing and allowing. The Philosophical Review. 1989;98(3):287-312.
19. Sulmasy DP. Killing and allowing to die: Another look. The Journal of Law, Medicine & Ethics. 1998;26(1):55-64.
20. Wolf SM. Chapter 10: Gender, feminism and death: physician-assisted suicide and euthanasia. In: Feminism and Bioethics: Beyond Reproduction. New York: Oxford University Press; 1996. p.282.
21. Velleman DJ. Against the right to die. The Journal of Medicine and Philosophy. 1992;17(6):665-681.
22. Miller FG, Truog RD, Brock DW. Moral fictions and medical ethics. Bioethics. 2010;24(9):453-460.
23. Dworkin R, Nagel T, Nozick R, Rawls J, Scanlon T, Thomson JJ. The philosopher's brief. In: Battin MP, Rhodes R, Silvers A, eds. Physician-Assisted Suicide: Expanding the debate. New York: Routledge; 1998. p.431-442.
24. Hart HLA, Honoré T. Causation in the Law. 2nd ed. Oxford: Clarendon Press; 1985.
25. Stauch M. Causal authorship and the equality principle: a defence of the acts/omissions distinction in euthanasia. Journal of Medical Ethics. 2000;26(4):237-241.
26. Callahan D. When self-determination runs amok. The Hastings Center Report. 1992;22(2):52-55.
27. Paterson C. On "killing" versus "letting die" in clinical practice: mere sophistry with words? Journal of Nursing Law. 2000;6(4):25-44.
28. Miller FG, Truog RD. Chapter 1: Withdrawing life-sustaining treatment: allowing to die or causing death? In: Death, Dying, and Organ Transplantation: Reconstructing Medical Ethics at the End of Life. New York, NY: Oxford University Press; 2012. p.1-25.
29. McGee A. Does withdrawing life-sustaining treatment cause death or allow the patient to die? Medical Law Review. 2014;22(1):26-47.
30. McGee A. Acting to let someone die. Bioethics. 2015;29(2):74-81.
31. McMahan J. Killing, letting die, and withdrawing aid. In: Steinbock B, Norcross A, eds. Killing and Letting Die. 2nd ed. New York: Fordham University Press; 1994. p.389.
32. Bronner B. The total artificial heart and the dilemma of deactivation. Kennedy Institute of Ethics Journal. 2016;26(4):347-367.
33. Jensen SJ. Killing and letting die. In: Jensen SJ, ed. The Ethics of Organ Transplantation. Washington, DC: Catholic University of America Press; 2011. p.170-191.
34. Asscher J. The moral distinction between killing and letting die in medical cases. Bioethics. 2008;22(5):278-285.
35. Trammell R. Saving life and taking life. In: Steinbock B, Norcross A, eds. Killing and Letting Die. 2nd ed. New York: Fordham University Press; 1994. p.290-297.
36. Bishop JP. Euthanasia, efficiency, and the historical distinction between killing a patient and allowing a patient to die. Journal of Medical Ethics. 2006;32(4):220-224.
37. Brody H. Causing, intending, and assisting death. Journal of Clinical Ethics. 1993;4(2):112-117.
38. McIntyre A. Doctrine of double effect. Standford Encyclopedia of Philosophy. Zalta EN, ed. Standford: Metaphysics Research Lab; 2019.
39. Anscombe GEM. Intention. 2nd ed. Oxford: Basil Blackwood; 1963.
40. Frey RG. Some aspects to the doctrine of double effect. Canadian Journal of Philosophy. 1975;5(2):259-283.
41. Graham G. Doing something intentionally and moral responsibility. Canadian Journal of Philosophy. 1981;11(4):667-677.
42. Kuhse H. The Sancity-of-Life Doctrine in Medicine. A Critique. New York: Oxford University Press; 1987.
43. Brody B. Withdrawal of treatment versus killing of patients. In: Beauchamp TL, ed. Intending Death: The Ethics of Assisted Suicude and Euthanasia. New Jersey: Prentice Hall; 1996. p.90-103.
44. Bosshard G, Fischer S, van der Heide A, Miccinesi G, Faisst K. Intentionally hastening death by withholding or withdrawing treatment. The Middle European Journal of Medicine (Wiener Klinische Wochenschrift). 2006;118(11-12):322-326.
45. Fuchs T. The notion of "killing". Causality, intention, and motivation in active and passive euthanasia. Medicine, Health Care and Philosophy. 1998;1(3):245-253.
46. Cavanaugh TA. The intended/foreseen distinction's ethical relevance. Philosophical Papers. 1996;25(3):179-188.
47. Cavanaugh TA. Double effect and the ethical significance of distinct volitional states. Journal of Christian Bioethics. 1997;3:131-141.
48. Cavanaugh TA. Act evaluation, willing and double effect. Proceedings of the American Catholic Philosophical Association. 1998;71:243-253.
49. Biggar N. Aiming to Kill: The Ethics of Suicide and Euthanasia. London: Darton, Longman & Todd; 2004.
50. Quill TE, Dresser R, Brock DW. The rule of double effect – a critique of its role in end-of-life decision making. New England Journal of Medicine. 1997;337(17):1768-1771.
51. Isaacs TL. Moral theory and action theory, killing and letting die. American Philosophical Quarterly. 1995;32(4):355-368.
52. Meilaender G. The distinction between killing and allowing to die. Theological Studies. 1976;37(3):467-470.
53. Ramsey P. The Patient as Person: Explorations in Medical Ethics. New Haven: Yale University Press; 1970.
54. Bishop JP. The Anticipatory Corpse: Medicine, Power, and the Care of the Dying. Indiana: Notre Dame Press; 2011.
55. Huddle TS. Moral fiction or moral fact? The distinction between doing and allowing in medical ethics. Bioethics. 2013;27(5):257-262.
56. Orentlicher D. The Supreme Court and physician-assisted suicide - rejecting assisted suicide but embracing euthanasia. New England Journal of Medicine. 1997;337(17):1236-1239.

57. Billings AJ, Block SD. Slow euthanasia. Journal of Palliative Care. 1996;12(4):21-30.
58. Kagan S. The additive fallacy. Ethics. 1988;99(1):5-31.
59. Asscher J. Killing and letting die: the similarity criterion. Journal of Applied Philosophy. 2007;24(3):271-82.
60. Atkinson G. Killing and letting die: hidden value assumptions. Social Science and Medicine. 1983;17(23):1915-25.
61. Ursin LØ. Withholding and withdrawing life-sustaining treatment: Ethically equivalent? The American Journal of Bioethics. 2019;19(3):10-20.
62. Beauchamp TL. Justifying Physician-Assisted Deaths. In: LaFollette H, ed. Ethics in Practice. 5th ed. Oxford: Wiley-Blackwell; 2020. p.78-85.