5

Two Dogmas of Research Ethics

5.1 Is There a Dilemma at the Heart of Research with Humans?

The historical reluctance in research ethics to embrace or recognize a social imperative to carry out medical research grows out of the worry that such an imperative too easily overrides and overshadows the rights and interests of individuals. I argued in the previous chapter that this worry is well founded when such a social imperative is cast in terms of the corporate conception of the common good. In contrast, the egalitarian research imperative that I outlined in §4.7 is predicated on the idea that advancing the generic interests conception of the common good through research with human participants is not fundamentally inconsistent with respecting the rights and welfare of study participants. In fact, the view I defend goes further, holding that respect for the status of individuals as free and equal is an integral, enabling component of the research enterprise understood as a voluntary scheme of mutual cooperation aimed at producing an important public good.

Even if we accept that the social imperative to create a system of research that advances the common good is also an imperative to ensure that such a system represents a voluntary scheme of mutual cooperation among free and equal persons, doubt might remain as to whether medical research can operate on those terms. Put another way, even if it is possible to ground some social or political institutions in the generic interest view of the common good, and to organize them in ways that are consistent with its requirements, it does not necessarily follow that the research enterprise is such an institution. In particular, the way that research exposes participants to risks, and the way that research ethics evaluates whether or not risks are reasonable or acceptable, might pose special problems for the egalitarian research imperative.

For the Common Good. Alex John London, Oxford University Press. © Oxford University Press 2022.
DOI: 10.1093/oso/9780197534830.003.0005

As we saw in §2.2, Walsh McDermott thought that the rule of law and freedom from arbitrary interference could not be extended into the realm of research with human participants because of the "moral dilemma of clinical investigation" (1967, 40–41). A moral dilemma is a situation in which every option an agent faces violates or transgresses some important norm or value. Agents who face a moral dilemma have to make tragic choices in the sense that every option available to them results in doing or allowing something that is bad or wrong (Levi 1986). They cannot extricate themselves from such a situation without incurring a moral loss. Even if more recent commentators reject McDermott's claims about the scope of researcher discretion, many share the fundamental perception that "tragic choices [are] involved in designing a system for research on human subjects" (Menikoff and Richards 2006, 19).

In this chapter I show how some of the problematic commitments that I identified in §1.2 and chapter 2 create a conceptual ecosystem in which the proposition that there is a deep and ineliminable conflict at the heart of research with human participants appears to be analytic, a conceptual truth about the nature of research and research risk. In particular, I show that these problematic commitments are often shared by protagonists on opposite sides of prominent debates and that this obscures their role in structuring the problem being discussed and the options for resolving them that are seen as salient or feasible. This critical or deconstructive work is thus necessary to clear the requisite conceptual space for an alternative framework for risk assessment and management within research ethics. In the next chapter I present such a positive framework and demonstrate how research risks can be managed in a manner that is consistent with a principle of equal respect that satisfies the requirements of the egalitarian research imperative.

This chapter examines a series of arguments that purport to show that there is a moral dilemma at the heart of research with human participants. Examining these positions highlights the central role of two largely unquestioned dogmas of research ethics. The first is the claim that the ethical norms that govern this activity derive from role-related obligations of professionals. The second is that clinical research is an inherently utilitarian undertaking. These dogmas are supported by, and lend support to, a functional view of clinical medicine and medical research that effectively identifies these activities with a set of goals and reasons that direct the individual decision-maker to optimize two incompatible metrics: as a clinician the decision-maker is obligated to provide optimal care to the individual patient but as a researcher

the decision-maker is obligated to generate the information that will advance the medical interests of future patients. Together, these dogmas structure the conceptual ecosystem in which genuine (but ultimately manageable) tensions within research appear to pose a fundamental dilemma that calls for tragic choices.

I have tried to order these arguments from those that are more general and wider in scope to those that only apply to research with particular features. In §5.2 I examine the most philosophically general argument which holds that the fiduciary duties of clinicians are necessarily incompatible with the utilitarian goals of research. In §5.3 I argue that statements about the logical or conceptual incompatibility of the ends of research and medical practice show only that these are distinct activities and do not establish that they cannot be organized in a way that reconciles respect for individual interests with pursuit of the common good.

The remaining arguments rely on more contingent features of the research enterprise to generate a moral dilemma. Nevertheless, they share a number of assumptions in common, and it is important to highlight the role of those assumptions in these arguments. To do this, in §5.4 I present what I call the *template for the appeal to uncertainty*. The template provides the most general formulation of the claim that uncertainty about the relative therapeutic, prophylactic, or diagnostic merits of a set of interventions for a particular problem offers a way to reconcile respect for the interests of study participants with the generation of socially valuable information. Whether this argument is sound depends on how a number of key claims are spelled out in practice.

Stating this position in its most abstract form and highlighting the role of these key claims that must be further specified is important for two reasons. First, the positive view I elaborate in the next chapter includes a version of this appeal. So, it is important to establish that there are many ways in which this template can be filled out, some of which resist the objections that are discussed in this chapter. Second, it allows us to show how arguments to the effect that the position outlined in the template are unworkable presuppose very particular ways of filling in some of its key features.

In §5.5 we examine one of the earliest and most influential views that fills out the template for the appeal to uncertainty on terms that have come to dominate the literature. In particular, Charles Fried (1974) argues that if studies begin in the relevant state of uncertainty—given the perhaps unfortunate name "equipoise"—and if they are designed to disturb that state of

equipoise, then they can reconcile the individual clinician's duty of personal care and the researcher's obligation to generate valuable information. Fried fills out the template in terms that presuppose a particular conception of uncertainty and that locate that uncertainty in the judgment or in the head of the individual clinician-researcher. Within the conceptual ecosystem of orthodox research ethics this way of filling out the template is natural and intuitive. But I show in §5.6 that it is also doomed to failure. This approach produces self-defeating practices that neither generate sound scientific evidence nor safeguard and advance the interests of individuals.

The failure of Fried's view and the fact that it appears natural and intuitive within the conceptual ecosystem of orthodox research ethics encourages the appearance of an intractable dilemma that arises from practical features of particular studies. But this natural and intuitive view, and the conception of equipoise that it entails, is only one from among a much larger universe of possible views. In particular, where Fried embraces a relatively fragile conception of uncertainty that is located in the head of the individual clinician-researcher, Benjamin Freedman articulates an alternative under the heading of "clinical equipoise" that locates the relevant uncertainty in the expert medical community, and that recognizes that uncertainty can arise from the conflicting assessments of experts who are not themselves uncertain about the merits of the interventions in question.

The fact that these views are often confused in the literature illustrates how deeply ingrained the two dogmas of research ethics are within the conceptual ecosystem of orthodox research ethics. Moreover, I show in §5.7 how the force of those dogmas has led even staunch proponents of clinical equipoise to question its moral relevance and to supplement that view with requirements that effectively recapitulate the problems associated with Fried's view. The upshot of these arguments is to show that common and intuitive ways of completing the template for the appeal to uncertainty are unworkable, but that the intuitive force of these views is rooted in the two dogmas of research ethics that I ultimately argue we should reject.

In §5.8 we turn to an argument that is still narrower in scope than those discussed previously but that appears to be more straightforwardly successful. This argument holds that the dilemma at the heart of research follows from the fact that research often requires participants to undergo risky or burdensome procedures that are not offset by the prospect of direct benefit

to those same participants. This poses a special problem for any view that appeals to the template outlined in §5.4 (including Freedman's clinical equipoise) since few experts are likely to be uncertain about the fact that study-related procedures impose risks and burdens on participants that are not offset by the prospect of direct benefit to those same individuals.

More generally, however, this argument has been used to show that research participation is antithetical to the rational self-interest of individuals and that this conflict between the rational self-interest of individuals and the value of research to the community produces a coordination problem known as the prisoner's dilemma (Heyd 1996; Wertheimer 2010, 9). As a result, studies that contain such purely research-related procedures are supposed to be antithetical to both the clinician's fiduciary duty to patients and to the participant's own rational self-interest.

In §5.9, however, I argue that any moral standard that treats the risks and burdens of purely research-related study procedures as antithetical to the clinician's fiduciary duties would be so restrictive that it would prohibit a variety of ethically permissible practices in clinical medicine. Since clinical medicine is the domain in which the clinician's fiduciary duties should be most clearly exemplified, the arguments of this section show that research ethics retains a last vestige of unjustified medical paternalism.

I also argue that arguments purporting to show that research participation is a prisoner's dilemma rely on a conception of individual welfare that is excessively narrow and limited to individual health interests. I show that if such arguments were sound, they would not only apply to research participants, but to researchers. Once we recognize that the way health interests factor into a person's life plan can differ across individuals, the claim that research poses a prisoner's dilemma is undermined.

Ultimately, this long chapter concludes with reasons to reject both dogmas of research ethics and the way that they create a conceptual ecosystem in which several types of morally relevant diversity are obscured. The first is diversity in the expert medical community regarding scientific and medical questions. The second is diversity in democratic societies regarding the life plans that individuals adopt and pursue and the way those diverse life plans shape individual attitudes toward various risks and benefits. These forms of diversity are morally relevant, in part, because a requirement of justice in a decent society is to create social space in which individuals have the real freedom to pursue a life plan of their own. It is precisely this diversity in

first-order life plans that makes it possible to satisfy the egalitarian research imperative.

In the following chapter I articulate the integrative approach to research risk. Like some of these early views, it holds that credible uncertainty has a special role to play in research ethics: ensuring that research with humans has scientific and social value and reconciling research participation with equal respect for the rights and welfare of study participants. Unlike those views, however, it rejects both dogmas of research ethics. As a result, it does not frame the central problem as reconciling the moral duties of conflicting social roles, and so the solution that it provides is not constrained by assumptions that are built into the traditional way of framing the problem. Readers who are primarily interested in my positive view can turn directly to that chapter.

5.2 Incompatible Ends?

5.2.1 The First Dogma: Moral Norms from Role-Related Obligations

The idea that there is a dilemma at the heart of medical research is bound up with two dogmas of research ethics. The first dogma is that the relevant ethical norms in this domain grow out of, and are grounded in, role-related obligations. Miller and Brody express this idea when they argue that in this domain "the basic goal and nature of the activity determines the ethical standards that ought to apply" to it (2003, 22 and 1998) and that the goals of clinical medicine and the goals of clinical research are "logically incompatible" (Brody and Miller 2003, 332). As a result, they argue, the dilemma at the heart of research ethics is a fundamental conflict between the incompatible demands placed on a single decision-maker by the moral duties of two conflicting social roles—that of the clinician and that of the researcher.

To understand the dilemma at the heart of research ethics, on this view, we need to understand the sense in which clinical medicine and clinical research are logically incompatible. This, in turn, involves seeing these activities as structured by different frameworks of reasons that can diverge in both principle and in practice. Since the social roles in question are roles for a single agent, if the reasons that structure them cannot be mutually satisfied, then research ethics will necessitate tragic choices between a set of basic and irreconcilable values.

5.2.2 Hippocratic Obligations: Patient-Centered Consequentialism

Within research ethics, the role of the clinician tends to be understood and explicated in very traditional, Hippocratic terms. For example, the *Belmont Report* provides a standard expression of the physician's duty of personal care when it says that "the Hippocratic Oath requires physicians to benefit their patients 'according to their best judgment' " (National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research 1979). The World Medical Association's 1964 *Declaration of Helsinki* holds that, "The Declaration of Geneva of The World Medical Association binds the doctor with the words: 'The health of my patient will be my first consideration' " (1964). The idea that the health of the patient must be the researcher's first concern was made more explicit in subsequent versions of the Declaration. For instance, the version from 2000 says, "In medical research on human subjects, considerations related to the well-being of the human subject should take precedence over the interests of science and society" (2000). The fundamental moral duty of the clinician is thus defined by the therapeutic obligation (Hill 1963; Fried 1974; Peto et al. 1976; Peto and Baigent 1998; Sackett 2000; Miller and Weijer 2006), sometimes called the "principle of therapeutic beneficence." The underlying idea is that "physicians should promote the medical best interests of patients by offering optimal medical care; and the risks of prescribed treatments are justified by the potential therapeutic benefits to patients" (Miller and Brody 2002, 4).

This traditional view of the provider-patient relationship has a relatively clear structure to it. There are two principal parties, the clinician and the patient. The patient relies on the clinician's expert knowledge and skill to advance the patient's medical best interests. In return, the clinician has a fiduciary duty to use his or her best medical judgment to advance the interests of the patient. When deciding whether or not to conduct a procedure or offer a test, the clinician thus has to consider the likely outcomes of that procedure and how they will affect the medical best interests of that patient. Other concerns are either irrelevant or have the status of secondary considerations that can play a role in decision-making only so long as they do not interfere with the morally primary goal of advancing the patient's medical best interests.

The Hippocratic conception of the clinician-patient relationship thus has the structure of a patient-centered consequentialism. It is a form of

consequentialism because the right act for the clinician to perform is determined solely by the goodness of the outcomes it is likely to produce. Like other forms of consequentialism, Hippocratic patient-centered consequentialism is grounded in the value of beneficence—the main moral consideration used to evaluate acts is their likely impact on the good of those affected. Like other forms of consequentialism, it also involves an optimizing conception of rationality. The clinician's duty is to choose the optimal act—the one that brings about the best consequences. However, unlike other forms of consequentialism, which tend to evaluate the consequences of acts in terms of their outcomes for all affected parties, impartially considered, Hippocratic ethics is patient-centered. This means that the consequences that matter when evaluating actions are limited to their impact on the individual patient. Similarly, whereas most forms of consequentialism are concerned with the goodness of outcomes in a very broad sense of the good, Hippocratic ethics is focused on the health or medical best interests of patients.

Thinking of the clinician-patient relationship in these terms dovetails nicely with the idea that clinicians have a special, fiduciary relationship with patients. In a fiduciary relationship, the clinician has a special moral duty to put the interests of the patient above all other concerns—including their own private and professional interests. The ground for this duty traditionally hinges on several factors. Clinicians have expert knowledge and skills that patients lack but which patients rely on to advance their medical interests. This creates an asymmetry in knowledge and power between the two parties. By entering relationships with clinicians, patients become dependent on clinicians in a morally special respect—they rely on the expert knowledge and skill of clinicians to safeguard and advance their medical interests without necessarily having the ability to independently assess and monitor the actions of the clinician to make sure that they are aligned with the patient's best interest. Asymmetric knowledge and power create a relationship of dependence fraught with the potential for domination and abuse. Treating the clinician patient relationship as fiduciary in nature helps to facilitate social trust by articulating clear expectations about the relationship between patient interests and competing concerns. The social enforcement of these expectations provides public assurance that breaches of that trust will not be tolerated (Miller and Weijer 2006).

Hippocratic patient-centered consequentialism internalizes the fiduciary nature of the social relationship between clinicians and patients into the morality of medicine itself. It erects the health interests of the patient as

the good to be optimized and it places the physician under a duty to use her best medical judgment to always choose the act—the intervention or course of care—that is most likely to bring about the best medical outcome for that individual.

When deliberating about how to manage the potential therapeutic advantages of an intervention given its possible adverse effects, the clinician has a moral duty to choose the course of care in which potential burdens and risks of care for a patient are offset by the prospect for therapeutic advantage for that same patient. As a result, "when physicians of integrity practice medicine, physicians' and patients' interests converge. The patient desires to regain or maintain health to relieve suffering; the physician is dedicated to providing the medical help that the patient needs" (Miller and Brody 2003).

When Miller and Brody say that the ends of clinical medicine and the ends of research are logically incompatible, they are asserting that these activities are structured by different frameworks of reasons that can diverge in both principle and in practice. If the defining goals of clinical medicine involve advancing the health interests of the individual patient, then there are no circumstances in which the reasons that are internal to clinical medicine should ground conduct inconsistent with the medical best interests of patients. The goals of clinical medicine and the interests of patients are aligned, in this view, because the framework of reasons that structure that activity necessarily tracks patient interests.

5.2.3 The Second Dogma: Research as Inherently Utilitarian

In contrast, "clinical research is dedicated primarily to promoting the medical good of future patients by means of scientific knowledge derived from experimentation with current research participants—a frankly utilitarian purpose" (Miller and Brody 2003, 21 see also 2007, 162). The claim that research with human participants is an inherently utilitarian undertaking is a second dogma of research ethics.[1] One reason for its status as a pervasive and often unquestioned assumption is that it appears to be analytic—a

[1] Miller and Brody here give voice to a set of ideas that is often expressed in different terms. For example, it was common in earlier discussions to speak more explicitly of the "problem of experimentation" as setting the terms on which it is permissible to take some lives in order to save more lives (Calabresi 1969) or in which the interests of some must be traded off against the interests of others (Fried 1974).

conceptual truth derived from reflection on the point and purpose of the research activity. If the goal of research is to generate the knowledge necessary to advance the medical interests of large numbers of future patients and if the goals of this activity define the norms that govern it, then researchers have a duty to act so as to generate the knowledge that will bring about these advances in future medical care. Without any clear check or constraint on the methods that researchers can use to promote this end, this position is treated as permitting trade-offs between the welfare of study participants and future beneficiaries of research.[2]

Treating the role-related obligations of clinicians and researchers as different forms of consequentialism sharpens the distinction between these activities in a way that makes them appear "logically incompatible." Whereas the ethical duties of the clinician have the form of patient-centered consequentialism, the ethical duties of the researcher have the form of an impartial, utilitarian consequentialism. As forms of consequentialism, both of these moral frameworks share a slightly narrower focus on health-related outcomes. Both also presuppose an optimizing rationality grounded in beneficence, directed at evaluating the rightness of individual acts by assessing the consequences those acts are expected to bring about. They diverge, however, in their accounts of whose interests matter when it comes to evaluating those consequences: the interest of the individual patient alone or the interests of all future patients who stand to benefit from improvements in the standard of care.

Because Hippocratic, patient-centered consequentialism focuses solely on the medical interests of the individual patient, the expert decision-maker is faced with a problem of comparing the relative value of different health states for the same individual. This is a kind of *intrapersonal* comparison of utility: will the burdens, harms, or risks associated with treatment A be outweighed or offset by sufficient benefits to make the provision of A superior to the provision of treatment B, given its burdens, harms, or risks and the offsetting benefits that might result to the patient?

In contrast, utilitarianism requires that the decision-maker go further and compare the value of outcomes across different individuals. These *interpersonal* comparisons traditionally involve summing the value or disvalue that

[2] Strictly speaking, from the narrow claim that the production of socially valuable information is a necessary condition of ethically permissible research, if follows only that research that lacks social value is morally impermissible. Nothing follows about the extent of the demands that can be placed on the interests of free and equal persons in pursuit of this goal. I return to this point near the end of the present chapter.

results for different individuals from different courses of action (Sen 1979). As a result, the considerations that determine whether to perform a test or to administer an intervention to one set of people include the likelihood that doing so will generate information necessary to improve the standard of care that is available to a different set of future people. Moreover, if research is a utilitarian enterprise and if performing procedures or providing interventions that expose study participants to serious harms or risks is necessary to bring about a sufficiently significant benefit to a large enough group of future people, there are no grounds internal to the research enterprise itself on which to block or prevent such sacrifices. As a result, in this view, there is no in-principle alignment between the interests of study participants and the framework of reasons that structure the research activity.

5.2.4 Reasonable Risk: Trading Risk to Some for Benefits to Others

That research is an inherently utilitarian undertaking seems to be reflected in the way that reasonable risks are defined in the field:

Definition of reasonable risk: Risks to subjects that are not offset by the prospect of direct benefit to the participant must be reasonable "in relation to the importance of the knowledge that may reasonably be expected to result" from the study (45 CFR 46.111[2]).

A trial can pose an acceptable degree of risk to participants even if those risks are not offset by the prospect of direct medical benefit to participants themselves. Rather, such risks can be justified if they are offset by the prospect that they are necessary to generate sufficiently valuable information. This seems to countenance the permissibility of trading risk of harm to a small group of study participants if it will purchase sufficient social benefit for others.

This conceptual analysis outlines the conceptual ecosystem within which disputes play out over how to reconcile this fundamental tension. It sets the terms in which debates are framed, and the interlocking claims that go into this formulation of the problem constitute assumptions common to otherwise warring camps. For instance, disputes about how to respond to this tension often take place against a shared framing of the problem as a conflict in the vantage point of a single decision-maker. Normally the decision-maker

in question is the medical professional who cannot simultaneously satisfy the demands of these competing and incompatible forms of consequentialism.

Even when they disagree about how to respond to this problem, competing sides often assume that the problem arises because the individual decision-maker has a duty to do what in her best judgment will bring about the best outcome. As a result, the idea that the central tension in research is a conflict in the objectives to be advanced by a single, rational optimizer is baked into the problem from the start. Against this backdrop the conflict hinges on the different metrics this individual decision-maker is required to optimize to bring about the best outcome—the medical best interests of the present patient or the medical interests of a large group of future patients.

5.3 No Easy Analytic Answers

5.3.1 Two Senses of Incompleteness

The claim that there is necessarily a moral dilemma at the heart of the research enterprise appears to represent a deep philosophical truth that follows from a conceptual analysis of the role of clinician and the role of researcher. Against the backdrop of the first dogma of research ethics, this focus on social roles makes sense because the moral norms that govern this sphere are taken to derive from role-related obligations. Each of these social roles pursues a logically distinct set of ends which are part of distinct systems of norms and obligations. The moral obligations of the clinician represent a form of patient-centered consequentialism while the obligations of researchers represent a form of impartial utilitarianism. Against the assumption that these frameworks are to be implemented by the same individual decision-maker, it looks like such a person would necessarily face a choice between optimizing two different metrics: fidelity to the interests of the patient before them and fidelity to science and the greater good.

Even if we assume for a moment that this argument is sound, what does it show? The main point I want to make here is that, although it establishes that these are conceptually distinct activities that advance different ends, it does not show that these activities cannot be integrated in practice in a way that respects the rights and welfare of study participants while generating socially valuable information. In part, this is because professional norms are incomplete in two ways: their guidance may not always be adequate in the face of

uncertainty and their guidance may not reflect broader considerations that fall outside the narrow confines of issues recognized by professional roles. Finally, conceptual arguments about the nature of professional roles are often insufficient to answer substantive moral questions because professional roles can be defined in myriad ways, each of which incorporates different responsibilities.

To make these arguments, it is helpful to make explicit an idea that tacitly motivates the conceptual analysis offered in the previous section and that has deep roots in Western philosophy. This is the idea that professions, such as medicine, are distinct bodies of craft knowledge, each of which can be defined by the distinct end that it pursues. For ancient Greeks, craft knowledge or *techne* is the paradigm of a body of knowledge, covering a discrete domain, geared to bringing about or producing a discrete set of ends or outcomes. Different forms of craft knowledge are defined by the pursuit of different ends: blacksmiths make implements from metal, carpenters make objects from wood, generals understand strategy and how to use troops and tactics to achieve victory. Similarly, medicine has a long history of being conceived of as a craft whose purpose is to benefit the patient through the production of health.[3]

Whether the guidance provided by such bodies of technical knowledge is authoritative depends on two kinds of incompleteness. The first concerns whether it has sufficient knowledge to reliably produce the well-defined products or outcomes that define them. Even when it is clear what properties an object or outcome is supposed to have, the guidance of such a body of knowledge becomes less authoritative as its ability to reliably produce that product decreases. The second concerns the degree to which one craft relies on some other body of knowledge to determine what properties its products ought to have in order to serve the larger purposes and ends of the user.

Although different bodies of technical knowledge are distinct and can therefore make competing demands on the same individual, they can also be mutually supportive in actual practice. The reason, as Aristotle was well aware, is that no narrow branch of professional knowledge has as its subject overall individual flourishing. Rather, each has as its defining end the production of some relatively narrow good—health, wealth, victory, and so on. But the question of how to make a good life out of those goods is not a

[3] In the opening of the *Republic*, Plato has a protracted discussion of medicine as a craft distinct from the craft of money-making. For its continued relevance to today, see London 2000a and 2020.

technical question. It falls into the domain of ethics and what Aristotle calls *phronesis* or practical wisdom, which, at the social level, is the domain of political philosophy. As a result, the all-things-considered judgments that we make about the limits on professional powers and prerogatives and the constraints on their conduct must be informed by a larger conception of the way that the goals and activities of various professions fit into a social order that reflects the fundamental value of individuals and their interests in making momentous decisions for themselves and in forging and pursuing a good life.

5.3.2 The Incompleteness of Medicine

Medicine is incomplete in both of these respects. First, it is often not clear how to safeguard or advance the health of a patient. For example, we may not understand the pathophysiology of a novel disease and there may not be direct evidence about the effects of various interventions on that disease. At best we may have a range of hypotheses about the mechanisms through which the disease attacks the body and about which possible interventions might represent the best way to bring about a clinical benefit in patients with this disease.

In situations in which it is not clear how to advance a patient's medical best interests, the guidance of individual experts becomes less authoritative. The reason is that the warrant for the claim that some act or course of care is obligatory is grounded, ultimately, in the prospect that it will actually benefit the patient. As it becomes uncertain whether patients are better off receiving one form of treatment for a particular medical condition rather than another (for example, intervention A or B), then randomizing that patient to receive A or B has the advantage of generating reliable medical evidence without knowingly compromising the health and welfare of study participants. We will consider this argument in more detail in a moment (§5.4).

Second, medicine is also incomplete in the second sense outlined previously. Even if we assume that the clinician's moral duties are appropriately modeled as a kind of patient-centered consequentialism, it is simply false that choosing an act that is less than optimal from this standpoint is the kind of wrong that creates a moral dilemma in this space. The reason is that the dilemma in question concerns the rights and interests of study participants

and it is perfectly reasonable and ethically permissible for patients or study participants to make choices that are not strictly optimal from the standpoint of the individual Hippocratic clinician. Moreover, the case for this claim is most compelling in precisely those circumstances in which the moral case for conducting research is the most compelling.

Consider the following example (London 2020). At the inception of the SARS-CoV-2 outbreak there was considerable uncertainty about the pathophysiology of this novel disease and about the best methods for preventing its spread and for treating infected patients. Experts relying on hypotheses about disease pathophysiology and about intervention mechanism constructed a list of at least a dozen interventions they regarded as likely to produce a therapeutic effect in patients. It included prednisone, dexamethasone, baricitinib, methylprednisolone, enoxaparin, colchicine, remdesivir, favipiravir, ivermectin, tocilizumab, lopinavir/ritonavir, azithromycin chloroquine/hydroxychloroquine, and convalescent plasma (Herper and Riglin 2020).

Imagine that for each of these *n* interventions there was a passionate group of clinicians who, looking at the largely indirect evidence that was available, was convinced that their favored intervention was likely to produce the best outcomes for patients. If we assume also that the morality internal to the role of caregiver requires that each recommend for their patients what each believes is likely to maximize the patient's health interests, then it would be impermissible for such researchers to recommend anything but their favored intervention to patients. This means that it would violate their Hippocratic duty to recommend participation in a clinical trial and it is difficult to see how they could refer a patient for a second opinion if they know that their colleagues prefer different treatments as likely to be best.

Paradoxically, however, it would be permissible for each patient to seek a second, or a third, or an *n*th opinion. Imagine, then, that some patients visit each of these *n* groups of experts who each favor a different intervention as likely to be medically best for this patient. It is permissible for each of these patients to decide which clinician they want to care for them, even though doing so results in a choice that *n*-1 experts regard as suboptimal. In other words, if a patient agrees to be treated by a clinician who recommends one of these interventions, then all of the others might regard this as a bad choice. But it is not wrong to permit patients in this situation from making such choice.

So here we have a case in which different clinicians recommend different treatments to a patient as likely to be best. If it is permissible for each to make such a recommendation, then it is not wrong for each to act in a way that *n*-1 experts regard as likely to bring about less than the best outcome. Similarly, the patient chooses an option that *n*-1 clinicians regard as violating their Hippocratic duty to do what is in the best interests of the patient, but this choice is not morally wrong.

As we will see in this and in the next chapter, if it is permissible for patients to choose at random which experts should provide their care, then it should also be permissible for those same individuals to choose the option of participating in a well-designed trial in which they would be randomized to one of these *n* interventions. For now, my point is simply that the abstract argument from the previous section fails to capture the two important respects in which medicine is incomplete. As a body of knowledge about how to produce health, it is incomplete in the sense that there will arise cases in which there is uncertainty or conflicting expert judgment about how to best advance the medical interests of patients. In those cases, research provides a way to generate this knowledge and, as I will argue in more detail in the next chapter, this can be done without compromising the rights or welfare of study participants. These activities may be conceptually distinct, but not only may their ends not be incompatible, but in order to fulfill its mission of translating therapeutic intent into actual benefits for patients, medicine may require the thoughtful conduct of well-designed research (London 2020).

Similarly, medicine is incomplete in the sense that the goal that it produces is not the highest good there is. Health is an important good, but its value relative to other ends is a question that falls outside of the technical bounds of medicine. Even if clinicians are bound by Hippocratic duties to always act in what they regard as the patient's best interests, patients are morally permitted to act in ways that subordinate their narrow medical interests to the pursuit of other goals and commitments. In the example just discussed, this takes the form of deciding to allow themselves to be randomized to 1 of *n* alternative treatments for their medical condition rather than deciding at random to receive care from one or another expert clinicians. In this case, participating in research advances an important social good without necessarily requiring a sacrifice of self-interest on the part of the participant. We will revisit this point several times in the remainder of this chapter.

5.3.3 The Incompleteness of Research

It is worth noting that research, as a technical body of knowledge, is also incomplete in this second sense. If we think of research as something like the craft whose domain is the scientific and statistical methods needed to generate reliable knowledge about the safety and efficacy of various kinds of interventions, then this craft is incomplete to the extent that it relies on other disciplines—such as health policy informed by a proper concern for the freedom and equality of community members—to articulate what knowledge gaps ought to be addressed and to articulate the constraints on permissible methods of addressing those gaps.

When it comes to determining what the constraints are on permissible research studies, the kind of conceptual analysis described in §5.2.1 is not sufficient to answer this question. It is helpful to see that this point can be put in two different ways. In both cases, even if we assume that the purpose of research is to generate knowledge that will advance the interests of future patients, it does not follow that such research is inconsistent with respect for the rights and welfare of study participants.

The first way to make this point is that if we follow the second dogma of research ethics and we grant that research is an inherently utilitarian undertaking, it does not follow that it is wrong, all things considered, not to conduct studies that are regarded as optimal from this narrow viewpoint. Studies that optimize social value may be morally wrong, all things considered, if they do so by abrogating the rights and interests of study participants. Many of the studies described by Beecher fall into this category (§2.2.3). Likewise, studies that fall short of optimality when narrowly considered may be ethically preferable to studies that are optimal in the narrow sense, if they generate sufficient social value to improve the capacity of social institutions to meet the needs of community members without violating or diminishing the rights or welfare of study participants in the process.

In fact, if it is a conceptual truth that research is in some sense an inherently utilitarian activity, then we might also say that it is axiomatic in research ethics as a field that this utilitarianism must be constrained. Asserting the conceptual incompatibility of the norms of research and any other set of norms simply amounts to saying that when research operates under such constraints it may not be optimally utilitarian. But so what. The question is whether, from the standpoint of a just society, it can produce the information necessary to improve the capacity of basic social structures to meet, secure,

and advance the basic interests of its various members on terms that respect the status of individuals as free and equal. No purely conceptual argument about the proper ends of this activity can establish that this cannot be done in actual practice.

5.3.4 If External Constraints Are Unnecessary the Second Dogma Is False

Alternatively, research can be defined in various ways and some have argued that the norms for limiting the demands that research can place on study participants can themselves be derived from features internal to the research enterprise. For example, it has been argued that, unlike physicians, who have a fiduciary duty to their patients, researchers have only the weaker obligation not to exploit study participants (Miller and Brody 2002, 2003). If this duty of non-exploitation is internal to research, then research isn't a fundamentally utilitarian undertaking after all and the second dogma of research ethics is false.

Research isn't a fundamentally utilitarian activity, on this assumption, because utilitarianism recognizes only a single duty—to perform the act that brings about the greatest good (see §3.4.3). But if researchers are forbidden from bringing about some real benefit if it involves exploiting study participants, then research would not be utilitarian. Among its goals and ends there would be a set of considerations of sufficient moral import that they sometimes outweigh the production of information necessary to improve the medical care of large numbers of future people. Views that contain considerations of sufficient weight to outweigh the production of greater good are not consequentialist and so cannot be utilitarian.[4]

Now the question arises as to why we ought to adopt this particular definition of research. After all, we can think of at least three conceptions of the research enterprise. Call the first "research," which is defined as a body of knowledge with the purely utilitarian end of maximizing the knowledge necessary to advance the standard of care for future patients. Call the second

[4] Such views are not consequentialist because consequentialism is the view according to which the goodness of outcomes is the only factor that determines the rightness or wrongness of an act (Kagan 1998). Views that accept that the goodness of outcomes matter, but hold that there are additional constraints on which actions are right or wrong, are forms of moderate deontology because they recognize constraints of sufficient strength that they sometimes outweigh the production of good outcomes.

"research*," which is defined as the production of such knowledge within the constraints imposed on that activity by a duty of personal care. Call the third "research**," which is defined as the production of such knowledge within the constraints imposed on that activity by the duty of non-exploitation.

By now it should be clear that for any package of constraints [x, y, z,] we can define a conception of "research***" that pursues those constraints by definition. In the face of competing definitions of competing practices, it is a substantive ethical and political question whether we ought to permit the conduct of research, since the demands that it can place on participants are not constrained by anything other than the prospect of helping future people. Perhaps, instead, we ought to forbid the practice of research and only allow the conduct of research*, since that produces socially valuable information and forbids the violation of a fiduciary duty to participants. Or perhaps we should forbid the practice of both research and research* and allow only the practice of research** or research***.

The point is that conceptual analysis can help us differentiate research from research* or research**, and so on, but it cannot settle the substantive moral question concerning which of these practices we ought to promote and how we ought to design the institutions that promote them. Substantive moral questions of this type cannot be derived from analytic claims, since such claims merely tell us how to define our words and concepts. Even if research and medicine are distinct bodies of technical knowledge, how their respective ends should be pursued in a just society and how their pursuit should be reconciled with ends of other activities and the needs of community members are substantive questions of ethics and policy that fall outside of the parochial expertise of either set of professionals.

The upshot of the argument so far is that the conceptual argument for the logical incompatibility of medicine and research cannot ground any substantive claims about how to tackle the challenge of integrating the potentially competing demands of these different disciplines within a just social order. At best, this argument shows that these undertakings are distinct, guided by different ends and responsive to different reasons. The norms that are internal to these activities and that are grounded in an understanding of the ends they pursue are technical norms about how to effectively apply these bodies of knowledge to bring about ends of a particular sort. All distinct bodies of technical knowledge are governed by distinctive norms of this type. But this does not pose an all-things-considered moral dilemma at a level of ethics or policy since the narrow, technical norms of such disciplines do not

extend into the larger ethical and political domain of how to integrate various activities within a just social order. Rather, what the scope of these productive disciplines should be, when to call on the one rather than the other, and how to reconcile their pursuit in a just society are substantive ethical and political questions that requires a broader set of values and concerns than the narrow technical norms internal to these disciplines.

5.4 Reconciliation through Uncertainty: The Template

The argument of the previous section rebuts the claim that from the fact that clinical medicine and research are conceptually distinct we can show that there is an inherent moral dilemma at the heart of research ethics. Even if the arguments made earlier are correct, however, there may be more practical grounds for concern about our ability to reconcile substantive requirements pertaining to the welfare of individuals with the features studies require to generate scientifically sound and socially valuable knowledge.

To motivate these worries we need to return to one of the arguments I outlined in §5.3.2. I used that argument to show that the goals of research are not necessarily inconsistent with the interests of study participants. I will first present this argument in schematic form as a kind of template in the sense that a number of its key propositions must be specified in more detail in order for the content of the argument to be clear in operational detail.

Understanding what I refer to as the *template for the appeal to uncertainty* is important for two reasons. First, I show how the two dogmas of research ethics make one way of filling out this template seem natural and intuitive (§5.5). The problem is that the resulting view is unworkable and doomed to failure. Because this view is often seen as the only way to complete this argument, the fact that this common and intuitive way of completing the template is unworkable reinforces the perception that there is a fundamental dilemma at the heart of research ethics.

Second, it is important to understand that the intuitive and natural way of completing this template is unworkable because of flaws that derive from the two dogmas of research ethics and not from flaws inherent in the template itself. Establishing this point is essential, in part, because in chapter 6 I provide a way of operationalizing the template that avoids those problems and redeems the ambition of reconciling the pursuit of social value with respect for the welfare of individuals. Since these are both ways of completing the

same general template, it is important to recognize how these views differ so that we can avoid confusion.

To lay out the template for the appeal to uncertainty, it is helpful to begin by giving more precise definitions to two requirements that appear to be in conflict. I will define the first, the Social Value Requirement, in a way that reflects the content of the egalitarian research imperative:

Social Value Requirement: Research with human participants is only justified if it is reasonably expected to generate the knowledge necessary to develop interventions, policies, practices, or other advances that will enable a community's basic social structures (such as its health-related institutions) to more effectively, efficiently, or equitably safeguard and advance the basic interests of its constituent members.

The social value requirement states a necessary condition for ethically acceptable research with humans. It is seen as in conflict with the following requirement:

Concern for Welfare: It is impermissible to knowingly expose a person to interventions, practices, or procedures that are known or credibly believed to be worse than another available option.

One of the most enduring and important ideas about how to reconcile concern for welfare with the social value requirement appeals to the existence of credible uncertainty. The template for this argument, in its most general form, can be stated as follows:

Template for the Appeal to Uncertainty: When there is credible uncertainty about the relative merits of the set of interventions available for addressing an important health problem, it does not violate concern for the welfare of study participants to allow them to be allocated to an intervention from that set by a method (such as randomization) that facilitates the production of reliable medical evidence.

The idea is that when it is clear that such a state of uncertainty obtains, a trial designed to resolve that uncertainty—to bridge that knowledge gap—can generate socially valuable knowledge without requiring the denigration or abrogation of participant welfare.

Consider first the claim that studies designed to address such uncertainty are likely to have significant social value. If there is uncertainty about which of several treatment options is best for patients with a given condition, then research that provides the evidence necessary to vindicate the clinical merits of one alternative over the rest would have a strong chance of altering clinical practice in a way that renders care more effective and efficient. In fact, uncertainty about the merits of the interventions being tested seems to be a necessary condition for sound science since research is a tool for learning, and if the answers to the questions posed are already known, then there is nothing to learn.

Uncertainty about the relative merits of interventions being tested also seems to be a necessary condition for socially valuable research since resolving uncertainty of this kind enables various stakeholders to better discharge important moral or social responsibilities. Clinicians can prescribe optimal care to patients. Patients have greater assurance about the likely effects of various courses of care. Health systems can make a more efficient use of scarce resources by implementing the best therapeutic, prophylactic, or diagnostic options and eliminating less effective care or practices. Policymakers will know which courses of care to promote, and perhaps also which lines of research to foster and support and which to abandon or demote.

Now consider the proposition from the standpoint of participant welfare. If the relative clinical merits of a set of interventions are uncertain, then there are no credible grounds for treating one intervention as superior to the rest. In this case, being allocated to one intervention, rather than the others, does not involve knowingly providing that person with a level of care that is known to be worse than another available option. In this case, allowing a patient to be randomized to the alternatives in this set does not violate or contravene the clinician's duty of personal care. Problem solved!

Unfortunately, matters are not so simple. What I've outlined previously is the template for the argument that scholars who appeal to uncertainty want to make. Part of the problem, however, is that the template is ambiguous about a nexus of specific claims or views that are tightly connected. To give this nexus of views specific content is to fill out the details of a framework of moral assessment within which the appeal to uncertainty has substantial moral content. As those views are given more precise content, the credibility of the argument that results can be evaluated more precisely.

The following four questions capture the nexus of issues that must be specified in order for a framework involving this kind of appeal to uncertainty to have determinant content:

1. **Normative Basis:** What is the normative basis for focusing on uncertainty?
2. **Whose Uncertainty:** Whose uncertainty matters when contemplating these issues?
3. **Model of Uncertainty:** How is "uncertainty" to be understood and modeled?
4. **Epistemic Threshold:** What is the window that determines when the relevant uncertainty obtains and when it has been removed or disturbed?

Against the background of the two dogmas of research ethics, what appears to be the most intuitive and natural way of specifying these views results in a position that cannot support or redeem the ambitions of the template I have laid out. Rather than rejecting the claim that uncertainty plays an important role in bridging concern for individual welfare and social value, as some have, I argue that we should reject the background views that make those unworkable assumptions seem so natural and intuitive. In order for uncertainty to play a critical role integrating ethical and scientific aspects of research, we must dispense with the two dogmas of research ethics.

5.5 The View of Equipoise That Refuses to Die

5.5.1 The Normative Basis for Appealing to Uncertainty

Proponents of the principle of equipoise, like Charles Fried (1974), Benjamin Freedman (1987, 1990), Paul Miller and Charles Weijer (2006a, 2006b) and proponents of the uncertainty principle (Hill 1963; Peto et al. 1976; Peto and Baigent 1998; Sackett 2000), ground the normative basis for focusing on uncertainty in its ability to render research participation consistent with the clinician's duty of personal care, or the fiduciary obligation to provide optimal care to each individual patient. On this view, uncertainty provides the key for turning one dogma of research ethics against the other: the best way to limit the inherent utilitarianism of the research enterprise is to circumscribe

the obligations of one professional role within another. The demands that can be exacted from patients are limited to those that are consistent with the clinician's duty of personal care.

5.5.2 Whose Uncertainty Matters

Thinking of uncertainty as a bridge between the goals of science and the moral duties of the individual clinician or researcher entails a particular view of the second question that fills out the template for the appeal to uncertainty, regarding whose uncertainty matters. If the moral obligations of researchers are derived from the physician's duty of personal care, then it follows that the duty of personal care binds the individual physician in the clinical context. Each physician is charged with benefiting their individual patients according to their best judgment. As a result, this requires that the *individual clinician or researcher* must be uncertain about the relative net therapeutic merits of the available interventions in order to recommend that a patient enter into a clinical trial. As a result, Fried (1974) and others (Peto 1976; Chard and Lilford 1998) argue that the uncertainty must reside in the mind of the individual clinician or researcher. After all, individual clinicians or researchers have a special moral obligation to the individuals in their care and they must enroll participants in studies or perform study procedures on individual participants.

The individualistic nature of the provider-patient relationship and the duties of clinicians seems to require that the relevant uncertainty must be located in the mind of individual clinicians. As one proponent of the uncertainty principle puts it:

> An ethical physician must do what is best for his or her patients. She cannot participate in a controlled trial if she is certain that one arm is superior to the others and that some of her patients will receive an inferior treatment by participating in the trial. It does not matter whether her certainty is based on formal scientific studies, on personal experience, on anecdote, on tacit understanding, or rules of thumb. Whether her certainty is in accord with or diverges from the view of the medical community is irrelevant. Uncertainty is a moral prerequisite for a controlled study. If we know what we should do, we should do it, not study it. (Enkin 2000, 758)

Here, the focus on uncertainty in the mind of the individual clinician is combined with a relatively fragile epistemic threshold according to which uncertainty is the absence of even anecdotal reason to expect that at least one of the interventions under consideration is superior to the rest.

5.5.3 Modeling Uncertainty

With regard to the third question about how to understand uncertainty, the focus on the judgments of the individual clinician or researcher who must discharge a duty of personal care to each individual before her suggests that uncertainty should be understood as a subjective state of the individual decision-maker. The traditional, Hippocratic understanding of the duty of personal care models it as a duty to choose optimal care for each individual. As a form of patient-centered consequentialism, the right act is the one that is best for the individual in question. In a situation in which the effects of interventions are known with certainty, then the clinician's obligation is to choose the option that will produce the largest net benefit to the individual. When the effects of interventions are not known with certainty—when they are subject to some element of chance or when the information that we have about them is scant or unreliable—then standard theories of individual rationality hold that the best option is the one that has the highest expected value. Expected value is the product of two factors: the magnitude of the expected net benefit to the individual and the probability of that benefit being realized or obtaining in practice.

In order for the clinician's duty of personal care to permit participation in a study where an individual will be provided one intervention, chosen by a random process, from a set of several options, the clinician must believe that none of those options is likely to be better than the rest. Against the background of the traditional, Hippocratic conception of the therapeutic obligation, this state obtains when there is no difference in the expected value of the interventions in the set of treatment options. In this case, each of these interventions is an equal bet in prospect, meaning that each has the same expected value for the participant. This also seems like a fairly natural and straightforward way to interpret the concept of equipoise—the expected value of each intervention is such that the judgment of the expert is "equally poised" between them.

5.5.4 The Threshold of Uncertainty

At what point has uncertainty been disturbed such that providing a patient with a particular option or set of options from the available set would violate the clinician's duty of personal care? On the view we have been entertaining so far, uncertainty is disturbed as soon as the clinician regards one intervention as having a higher expected value than the rest. Once one intervention has a more favorable expected value than the other options, the clinician's duty of personal care is no longer indeterminate. Rather, their duty is to provide the option that they regard has the highest expected net benefit for the patient in question.

5.6 Doomed to Failure

5.6.1 The Fragility of Individual Uncertainty

One of the most damning objections to this very natural and intuitive way of understanding equipoise is that it is incapable of supporting or redeeming the project of reconciling concern for the welfare of study participants and the production of valuable scientific information. The source of this failure lies in its conception of individual uncertainty, which is so fragile and evanescent that it rarely obtains. As Marquis (1983) and others (Gifford 1986; Hellman 2002) argue, only in relatively rare circumstances will a physician believe that it is equally probable that two or more therapeutic options offer a particular patient the same degree of benefit. There will almost always be some bit of information or some aspect of one intervention that tips the balance of the clinician or researcher's subjective assessment in favor of one intervention over others. Because such a fragile state of uncertainty will rarely exist, clinical trials between therapeutic alternatives cannot ethically be initiated.

Alternatively, even if such a fragile state of uncertainty did obtain at the start of a trial, critics argue, it would vanish as soon as evidence from the trial emerges (Hellman 2002). As a result, equipoise will not persist long enough to bring a clinical trial to its desired conclusion. As soon as the trial generates its first data points the physician is obligated to update her beliefs about which intervention is most likely to best advance the patient's health interests. If one option appears to fare better than another, the hypothesis that one option is inferior to the other would be more probable than its complement. As

Marquis argues, this means that the clinician's therapeutic obligation is no longer neutral between available options since, "A physician should not recommend for a patient therapy such that, given present medical knowledge, the hypothesis that the particular therapy is inferior to some other therapy is more probable than the opposite hypothesis" (1983, 42). Once this fragile state of uncertainty is disturbed, the trial can no longer be justified.

Since the point of the appeal to uncertainty is to reconcile the production of socially valuable knowledge with respect for the welfare of study participants, the considerations laid out in this section are sufficient to show that the particular conception of equipoise we are discussing here is an abject failure. In the rest of this section I show that this view is subject to additional shortcomings. Before moving on to the additional problems with this view, however, it is important to understand exactly what the objections in this section show.

Because the conception of equipoise that refuses to die is often treated as synonymous with equipoise in general, some critics take the argument in this section to show more than it does. In particular, they take it to show that there is a moral dilemma in research in which we can either respect the individual welfare interests of study participants or we can generate the information necessary to promote scientific progress. We can't do both (Marquis 1983; Gifford 1986; Hellman 2002; Miller and Brody 2003). But this is a mistake.

The arguments in this section drive a stake into the heart of a very particular way of filling out the content of the template for the appeal to uncertainty. In particular, we have shown only that these objections apply to a very specific view, namely, the conception of equipoise in which the uncertainty in question resides in the head of the individual clinician or researcher and is represented as a subjective judgment that the interventions in question have equivalent expected therapeutic value. From the fact that this conception of equipoise prohibits the vast majority, if not the entirety of socially valuable clinical research, it does not follow that the template for the appeal to uncertainty is unworkable. That is because there are other ways that the template can be filled out. In particular, Freedman rejects the view of equipoise that locates the relevant uncertainty in the mind of the individual researcher, in part, to avoid these very objections. In the following chapter I argue for an alternative that is sufficient to reconcile respect for the welfare of participants with advancing the common good.

The point of these remarks is to highlight the significant influence of the various background claims and presuppositions that structure the

conceptual ecosystem in which this view of equipoise most naturally arises. If we accept all of those presuppositions, then not only does there appear to be a dilemma at the heart of clinical research, but it appears to be stark and bleak. On the one hand, if we endorse the idea that the best way to make sense of respect for the welfare interests of study participants is to appeal to the individual clinician's duty of personal care, and we retain the traditional, Hippocratic conception of that duty, then we have to bite the bullet and hold that most clinical research is unethical. Alternatively, if we step back from all of this, and examine our intuition that a good deal of sound clinical research is not morally objectionable, then we seem to have to bite a different bullet and infer the utilitarianism of clinical research has to take priority over concern for individual participant welfare. This is why it is critical to distinguish the template for the appeal to uncertainty, and the range of alternatives for completing its practical content, from this particular attempt to specify its content.

5.6.2 Permitting Senseless Studies

The previous section recapitulated some prominent arguments in the literature on equipoise. Those arguments show that the view of equipoise that refuses to die is overly restrictive in that it would prohibit scientifically and socially valuable research from ever starting and that, even if such research can be initiated, this framework would prevent it from generating sufficiently reliable information to alter clinical practice and advance the standard of care.

One potential weakness of that argument is that proponents of the conception of equipoise that refuses to die might bite the bullet and simply hold that it is not permissible to violate the clinician's duty of personal care as they conceive it. The important insight in this response is that the arguments in the previous section rely on an independent judgment that clinical research is sufficiently valuable and important that a view which would prohibit all research must be morally flawed. Without independent support for the value of research, the previous argument might be seen as begging the question—as asserting that research is morally acceptable when the argument in question gives us credible reasons to believe that it is not. As a result, it is important to consider other weaknesses in the view of equipoise in question that do not rely on any claims about the moral value of activities that such a view prohibits.

The next argument points out something that is not well understood in the larger literature: the view of equipoise that locates uncertainty in the mind of the individual clinician and treats it as a fragile state of equivalence in prospect is too permissive—the studies that most clearly satisfy its requirements can lack significant social value.

On the view of equipoise in question, research is permissible when the expert clinician or researcher does not have a preference between any members of a set of interventions. This condition occurs when the expert judges that the members of that set have equal expected value. If a clinician believes that interventions A and B have equal expected value, then, on this view, it is permissible to allow study participants to be randomized to receive either of these interventions. The objections of the previous section target the state of affairs in which we need new information to clarify the relative merits of a set of interventions in order to close a knowledge gap. I say they target this case because closing an information gap is a paradigm case of a study with social value—generating reliable information in that case has a high likelihood of altering clinical practice and providing patients and other stakeholders with the information they need to make momentous decisions. The point of the objections of the previous section is that the conception of equipoise that refuses to die is incapable of generating information of this kind.

But consider the case in which an agent's belief that the interventions in the relevant set are of equivalent clinical value rests on considerable prior evidence. In this case, subjecting these interventions to further testing would not have significant social value and so would not be a wise use of scarce resources since we are asserting, by hypothesis, that there is no evidence gap that needs to be filled to improve the care of future patients. Nevertheless, it would be the case that such a study is morally permissible on the view of equipoise under consideration since the individual clinician or researcher has no grounds on which to prefer one intervention over the others. Although there is no social value in initiating such a study, the conception of equipoise in question regards the study as morally permissible because it begins in a state of equipoise—the individual researcher regards the relevant interventions as an equal bet in prospect.

Moreover, this study represents one of the rare cases in which this view of equipoise can support a trial that runs to completion. In other words, if the interventions in question really are of equivalent value, then it would be possible to run the study to completion as long as interim evidence accurately reflects their equivalent practical utility. But the fact that such studies could

be run to completion cannot be used to rebut the present objection since the probability that such a study runs to completion is in inverse proportion to the social value of the study. The less likely it is that we will learn anything new, the more likely it is that such a study will remain permissible on this (faulty!) conception of equipoise.

The argument in this section does not rest on any question begging assumptions about the relative importance of generating socially valuable information compared to the importance of respect for participant welfare. Instead, it shows that the conception of uncertainty that motivates this common and intuitive view of equipoise does a poor job of tracking the social value requirement, since the clearest cases in which this type of equipoise is likely to obtain are the least likely to generate information that will close important information gaps.

5.6.3 Conflicting Judgments and Self-Defeating Requirements

There is an additional argument against the conception of equipoise as a fragile state of uncertainty in the mind of the individual clinician that does not rest on what the proponents of that view may be motivated to regard as potentially objectionable premises about how to trade off one value against another. In particular, this conception of equipoise is self-defeating in the sense that it sets back the legitimate interests of a range of stakeholders without advancing any countervailing interests. Part of the problem relates to the reason why Freedman rejected this view of equipoise, namely, it is incapable of addressing a common kind of medical uncertainty precisely because that uncertainty is not to be found in the head of any particular expert.

One of the reasons that Freedman rejected the view that equipoise is a state of uncertainty in the mind of the individual clinician is that he had a clear sense that this focus is too narrow. In particular, there is often reasonable diversity in judgment among well-qualified and informed medical experts. Any conception of uncertainty that focuses solely on the judgment of a single medical expert will fail to capture an important form of uncertainty that clinical research should play a key role in addressing. This is uncertainty in the form of conflicting expert judgments or conflicting medical assessments.

To illustrate this point, let us more carefully consider what states might count as examples of medical uncertainty. The view of equipoise that frames

uncertainty as a feature of the decision process of an individual decision-maker can recognize two states of affairs that count as exemplifying medical uncertainty.

I call the first state of affairs *clinical agnosticism*. This state obtains when the individual decision-maker does not have a considered expert judgment about the relative merits of a set of medical interventions for patients with a specific medical condition. For instance, we can imagine a condition for which there are no interventions that have been established as effective treatments and there is a novel intervention that has been shown to be safe in healthy adults. Now consider a proposal to test the efficacy of this novel intervention in a trial in which all participants receive usual medical care and randomization is used to determine which patients receive a placebo and which receive the novel, investigational intervention. In this case, the set of relevant interventions includes A, the investigational drug, and B, the placebo. An expert clinician is agnostic about the relevant merits of A and B if that expert considers the evidence so sparse or unreliable that it doesn't favor one option over the others.

Clinical agnosticism is different from what I will call *clinical equivalence*. This is because the agnostic clinician is unable to form a preference of any kind between the interventions in question. In other words, the expert is unwilling to say that A is preferable to B, that B is preferable to A, or that A and B are of equivalent value. In contrast, clinical equivalence is the state in which the expert believes that the evidence is sufficient to warrant a judgment about the relative merits of the interventions in question and, on the basis of that information, concludes that they are of equivalent expected value.

Clinical equivalence may itself come in different forms or flavors. For example, if the evidence about the relative merits of A and B is relatively sparse, then the assessments of the likely expected value of each intervention might involve probability distributions that are very wide and encompass a broad range of possibilities. In contrast, as evidence accumulates about the merits of A and B, that uncertainty might narrow, indicating a greater confidence on the part of the expert about what to expect from the provision of these interventions.

For our present purposes, the main point is that neither clinical agnosticism nor clinical equivalence can capture another state of affairs that seems to represent a paradigm case of medical uncertainty. I call this state *clinical conflict.* The state of clinical conflict obtains when one group of well-informed and expert clinicians have a strict preference for one option over

the others (for example, these clinicians regard A as a superior treatment option to B) but there are other, equally well-informed experts who regard B as the superior treatment option to A. This is a case of clinical conflict because every expert clinician has a definitive expert judgment that one intervention is superior to the others, but the judgments of these well-informed medical experts do not agree. In the example I've just given, no individual expert is in a state of clinical agnosticism or clinical equivalence with respect to A and B. Nevertheless, the community is in a state of clinical conflict because some judge A to be superior to B while others judge B to be superior to A.

Freedman rejects the idea that equipoise is a state of uncertainty in the mind of the individual clinician precisely because he recognizes that what I am calling clinical conflict is a form of medical uncertainty in which there is "a split in the clinical community, with some clinicians favoring A and others favoring B" (1987, 144). This is why he is at pains to say that his favored position, what he calls *clinical equipoise,* is "consistent with a decided treatment preference on the part of the investigators. They simply recognize that their less favored treatment is preferred by colleagues whom they consider to be responsible and competent" (1987, 144).

Freedman makes this move because he recognizes that there is significant social value in conducting research that has the prospect of reducing conflict among expert clinicians. If it is the case that one group of clinicians is correct, and that, for example, the clinical merits of A dominate the merits of B, then demonstrating this fact will reduce inefficiencies in current practice since, without the study, some experts would provide B to patients. Not only are some patients receiving inferior medical care, but scarce resources are being spent on the provision of inferior care. Reducing or eliminating such inefficiency will directly benefit patients and help health systems steward shared resources to more effective uses.

The problem, however, is that the view of equipoise that we have been entertaining here—the one that Freedman rejects and that refuses to die—cannot permit research that is designed to address a state of conflict among well-informed medical experts. The reason for this is simple: no expert is uncertain in the sense of uncertainty that defines that view. In a case of clinical conflict, every expert has a definitive expert judgment in favor of one option and no clinician is in a state of clinical agnosticism or clinical equivalence. Those who favor A over B thus see themselves as having a duty of personal care to their patients to provide them with A and to prevent them from being randomized to B. Similarly, those who favor B over A see themselves as

having a duty of personal care to their patients to provide them directly with B and to prevent them from being randomized to A. Because no clinician is uncertain, if the view of equipoise that refuses to die is correct, a trial that would establish the relative merits of A and B cannot be run.

This result, however, is absurd. Patients who happen to live in one place, or who happen to have a particular insurance provider, or who happen to be assigned to a particular clinician, will receive intervention A. Other patients, who happen to live in a different place, or who happen to have a different insurance provider, or who happen to be assigned to a different clinician, will continue to receive intervention B. Each clinician believes she is doing what is best for her individual patient, and each disagrees with the treatment recommendations of other equally well-informed medical experts. Prohibiting patients from being randomized to A or B results in a situation where arbitrary differences in location, insurance coverage, or other circumstances result in some patients receiving A and some receiving B, but under conditions in which the relevant merits of these interventions cannot be compared. The prohibition on randomization thus deprives a wide range of stakeholders of information that is relevant to decisions that affect people's health and welfare without advancing any countervailing interest.

Now consider the situation in which we allow patients to be randomized to A or B. In this case, some patients receive intervention A and some receive intervention B—just as in the status quo. Only now, randomization creates the conditions under which the effects of each intervention are statistically independent of a wide range of factors that might influence and confound the observed outcomes. As a result, the random allocation creates the conditions under which we can discern the relative clinical merits of A and B. We can learn, that is, whether one of these interventions is superior to the other.

Prohibiting this trial makes no one better off. It doesn't advance the interests of any person. Nor does it protect patients from receiving substandard care since prohibiting the trial permits both A and B to be provided by clinicians who favor them. All it does is deprive clinicians, patients, and other stakeholders of the information they need to better advance the health needs of people with this medical condition. Similarly, permitting this trial does not make anyone worse off. Each participant receives a level of care that would be recommended for them by an expert clinician. But when participants are matched with treatments by a random process, we can expeditiously learn about the relative merits of these interventions and improve medical practice.

Prohibiting a patient from being offered the option of being randomized to A or B in the presence of effective study oversight is self-defeating. It removes an option that does not adversely affect any interest of any stakeholder, but which does represent an avenue through which participants can contribute to the resolution of a clinically meaningful question. Blocking research as an avenue through which patients can contribute to the common good, in the sense defended here, is an unjustifiable restriction on individual liberty. Because it also prevents the generation of a social good without any offsetting benefit to participants, it stymies a socially beneficial undertaking without warrant. The myriad stakeholders who rely on the information such studies are intended to generate are deprived of that information, setting back the interests of the various stakeholders who depend on them, without an offsetting benefit.

The argument outlined here represents a powerful objection to the view of equipoise that requires uncertainty in the mind of the individual researcher. It identifies an area of uncertainty—clinical conflict—which that view is incapable of accommodating. Nothing in these arguments presupposes controversial claims about how to trade off risks to participants against the likely gains in socially valuable information. Rather, the kind of case outlined here represents a situation in which the interests of participants and the requirements of sound science are not in conflict. The fact that the conception of equipoise that refuses to die prohibits research in this case reveals the extent to which it is misguided, and its normative foundations fail to track the ethically relevant issues.

5.6.4 Confusion in the Field: The Uncertainty Principle, Equipoise, and Clinical Equipoise

Too often, the view that treats the relevant uncertainty as a fragile subjective state of the individual clinician is treated as synonymous with *the* equipoise requirement, or as capturing the essentials of all variants of the equipoise principle. For example, Ashcroft describes clinical equipoise as

> equipoise in the mind of the intending physician regarding treatment options. In many ways, this remains the best formulation. For clinical equipoise is a necessary condition on entering a patient into a trial, and if any clinician is not in clinical equipoise regarding a patient of a trial, then this

> (or any other of his patients) should not be entered by him or her into the trial. The ethical duty of the physician here is clear enough. (1999, 320)

Equating this position with clinical equipoise is a mistake. Nevertheless, it is a mistake whose seeds were sewn at the birth of the concept of clinical equipoise. When Freedman opens his seminal paper "Equipoise and the Ethics of Clinical Research," he writes:

> In the simplest model, testing a new treatment B on a defined population P for which the current accepted treatment is A, it is necessary that the clinical investigator be in a state of genuine uncertainty regarding the comparative merits of treatments A and B for population P. If a physician knows that these treatments are not equivalent, ethics requires that the superior treatment be recommended. (1987, 141)

In this general introductory statement, Freedman is following Charles Fried's formulation in which the uncertainty that is required to justify the trial is situated in the mind of the individual clinical investigator. This gives readers the false impression that Freedman is expressing his own, considered view in this passage. The problem is that Freedman does not endorse this view. He calls it "theoretical equipoise," which he rejects.

According to the view of clinical equipoise that Freedman actually endorses, the requisite uncertainty is located in the larger expert medical community. Equipoise obtains when "there is no consensus within the expert clinical community about the comparative merits of the alternatives to be tested" (1987, 144). Moreover, Freedman explicitly states that clinical equipoise can exist in situations in which no individual clinician is uncertain. This happens when there is "a split in the clinical community, with some clinicians favoring A and others favoring B." In this case, he argues, clinical equipoise is "consistent with a decided treatment preference on the part of the investigators. They simply recognize that their less favored treatment is preferred by colleagues whom they consider to be responsible and competent" (1987, 144). Finally, Freedman adopts a more robust epistemic threshold, according to which the relevant uncertainty persists until evidence for the superiority of one intervention emerges that would be sufficient to forge a consensus in the relevant expert clinical community (Freedman 1987, 1990). This threshold requires that the evidence supporting a claim to superiority on behalf of one intervention from among the set under consideration must

be sufficiently compelling that it will influence the practice behavior, not just of one physician, but of the community of physicians.

What Ashcroft identifies as "clinical equipoise," therefore, is actually what Freedman identified as "theoretical equipoise" and what Fried had referred to simply as "equipoise." Adding to the confusion, within literature from the United Kingdom this latter position (what Freedman calls "theoretical equipoise") is commonly referred to under the name of "the uncertainty principle" (Hill 1963; Peto et al. 1976; Peto and Baigent 1998; Sackett 2000). In contrast, what Freedman actually describes as "clinical equipoise," Ashcroft calls "collective or professional equipoise," a term that is also more common among writers from the United Kingdom (e.g., Chard and Lilford 1998).

As a result, when the concepts of equipoise or clinical equipoise are invoked in all but the most scrupulous literature, they are often glossed in the terms I outlined in §5.5—as a requirement that the individual researcher believe that the interventions in question are an equal bet in prospect. Even when scholars distinguish the concept of clinical equipoise from other variants of equipoise, the former view is frequently mislabeled. Given the proliferation of different nomenclatures, this has created a fair amount of both confusion and frustration.

As a result, the view I've outlined is like a character from a horror film. It can be shot, stabbed, and burned, but just when you divert your attention it rises again to stalk the pages of journals and lecture halls, reigning terror in its wake. In part, this happens when scholars who believe they have vanquished this view under one label—they have repudiated it under the label of "theoretical equipoise," for example—go on to invoke the content of theoretical equipoise in some more restricted domain (§5.7). In other cases, rampant confusion over what constitutes equipoise in general, or clinical equipoise in particular, promotes the tendency to (mistakenly) explicate any proposal made under this moniker by reverting to the terms of the view it is intended to displace. This process is undoubtedly fueled, in part, by the fact that the term "equipoise" seems to connote something like views that are "equally poised" on a scale or an edge of some sort.[5] This imagery, in a conceptual

[5] Eyal and Lipsitch (2017) is a recent example in which equipoise is rejected under the assumption that it requires individual uncertainty and an equal balance of probabilities.

ecosystem structured by the two dogmas of research ethics explicated previously, creates a set of entailments that seem entirely natural and straightforward. As a result, this conception of equipoise is the philosophical equivalent of the alien that has laid its egg in the stomach of its unwitting victim so that the monster can dramatically burst forth from the victim's chest, only the victim here is clinical equipoise and the view that bursts forth is the view of equipoise it was meant to supplant and replace.

For the moment, the point I want to drive home is that the view of equipoise that refuses to die—what Freedman calls "theoretical equipoise" and what Charles Fried simply called equipoise—cannot provide a workable foundation for scientifically and socially valuable research. This failure reinforces the perception that there is a fundamental conflict—a moral dilemma—at the heart of research ethics. As I showed in the previous section, Freedman had already recognized some of the weaknesses in this view and they drove his attempt to defend an alternative view that might avoid these shortcomings.

Freedman's core innovation was to move the uncertainty that is relevant to establishing the boundaries of morally permissible research out of the head of the individual clinician. I think that this move was largely correct, and I extend and build on it in the next chapter. However, because Freedman continued to ground his view in the role-related obligations of physicians, his view also suffers from significant problems.

Before turning to those arguments, it is important to consider a recent challenge to Freedman's conception of equipoise that has been articulated, perhaps surprisingly, by two of its most ardent defenders. Miller and Weijer (2006b) argue that clinical equipoise is insufficient as a moral safeguard on research because "clinical equipoise does not adequately specify the doctor-researcher's duty of care to the patient-subject" (2006b, 546). Examining their claims will underscore the extent to which a focus on the moral responsibilities of individual clinicians has such a powerful hold on the moral debate in this area. It also allows us to investigate the merits of an alternative formation of Fried's equipoise in which the relevant uncertainty is located in the mind of the individual clinician, but the fragile threshold for disturbing equipoise is replaced by a more robust, social threshold.

5.7 The Duty of Care Revisited

5.7.1 Does Clinical Equipoise Address the Wrong Issue?

Miller and Weijer frame the fundamental problem of research ethics in terms that recapitulate the first dogma of research ethics. On the one hand, randomized clinical trials (RCTs) are designed to produce the public good of generalizable medical evidence. On the other hand, physician-researchers owe patients a "duty of care" that requires that they exercise their discretionary powers to advance patient interests "to the greatest extent possible" (2006b, 545). They thus hold that the "central dilemma of the randomized clinical trial" arises "because offering patients enrolment in RCTs imperils the doctor's duty to act in their interests" (542). The core question to be resolved, then, is "when may physicians, consistent with their duty of care to patients, offer them enrolment in an RCT?" (542).[6]

Miller and Weijer are proponents of clinical equipoise, but they part ways with Freedman when they argue that "clinical equipoise does not adequately specify the doctor-researcher's duty of care to the patient-subject" (2006b, 546). Their argument for this claim involves several steps. First, they note, correctly, that questions about the social value of a trial and the reasonableness of the risks that it involves must be addressed at the point when a study protocol is being formulated and prior to the enrollment of study participants. In other words, before participants can be approached with the possibility of participating in a study, an IRB must find that the study is ethically permissible.

Second, Miller and Weijer argue that Freedman's clinical equipoise is the appropriate standard for approving a study protocol. In other words, IRBs can ask whether there is honest and informed disagreement among experts in the relevant medical community about the interventions to be tested in a study and, if this is the case, they can permit a trial to move forward. Miller and Weijer refer to this as fulfilling the state's obligation in protecting the "agent-neutral interests of patient-subjects" (2006, 543). Although they do not define the term "agent-neutral interests" the idea appears to be that these are interests that agents can be presumed to have insofar as they are patients with a particular medical condition who meet the conditions listed

[6] There is a minor typographical error in this passage that I have corrected in my quotation.

in the protocol's inclusion criteria and who lack the various characteristics listed as exclusion criteria. When IRBs find that clinical equipoise exists—that there is honest disagreement about the merits of the interventions in a trial for patients who meet the stated inclusion criteria and lack characteristics in the exclusion criteria—then the IRB ensures that participants "will not be asked to accept substandard treatment to participate in clinical research" (544).

Thirdly, they argue that because IRB approval is limited to the protection of these agent-neutral interests, such approval "does not entail the moral or legal acceptability of enrolling particular patient-subjects in research, nor does it entail the acceptability of their continued participation in the study, as these acts engage the agent-relative interests of patient-subjects" (2006, 545). Once again, the term "agent-relative" interests is not defined, but from the context it appears that it refers to specific or unique interests that pertain to individual subjects. Thus, for example, if the specific medical history of a patient suggests that receiving a particular intervention A would be "unduly harmful" (2006, 546) then it would be impermissible to enroll such a person in a study in which they might be randomized to A, even if that patient has the medical condition that A is intended to address.

As a result, Miller and Weijer argue that clinical equipoise captures a duty that the state owes to individuals who agree to participate in research to ensure that their agent-neutral interests will be protected in the course of such participation. However, as they understand it, "clinical equipoise does not contemplate the particular circumstances of individual patient-subjects. Therefore, it is not, and indeed cannot be, considered to be an adequate specification of the duty of care of doctor-researchers, because they are bound to protect the agent-relative welfare interests of the patient-subjects" (2006, 546). In effect, they argue that clinical equipoise is a solution to the wrong problem: the fundamental dilemma at the heart of research is about how to reconcile research participation with the clinician's duty of personal care—a subject, they argue, Freedman's clinical equipoise simply doesn't address.

5.7.2 The Clinical Judgment Principle

As a result, Miller and Weijer claim that although clinical equipoise is a necessary condition for ethically initiating a trial, no person can be enrolled into a trial solely on the basis of clinical equipoise. Rather, to reconcile study

participation with the physician's duty of personal care, it must also be the case that the individual physician-researcher regards study participation as consistent with that duty. And, like Fried and others, Miller and Weijer argue that the duty of care requires the exercise of discretionary powers for the sole purpose of advancing the individual patient's medical best interests.

Miller and Weijer take themselves to be showing that clinical equipoise and Fried's equipoise are not mutually exclusive. The former governs the review of study protocols by IRBs and the latter states the conditions under which individuals can be recruited into a study. How then do they propose to avoid the problems that led Freedman to reject Fried's view in the first place—the problems we canvassed in §5.5–5.6?

Miller and Weijer argue that it is a mistake to assume that the clinician's duty is based on a fragile epistemic threshold in which a mere hunch that one intervention is superior to the rest is sufficient to trigger the physician's duty of personal care and require the provision of that intervention and no other. Instead, they argue that individual researchers are subject to what they call the "clinical judgment principle," which holds that if an RCT has been approved by an IRB, "the physician may offer patients enrolment in a trial unless (1) they believe that it would be medically irresponsible to do so and (2) this belief is supported by evidence that ought to be convincing to colleagues" (2006, 546).

How is this clinical judgment principle supposed to avoid the problems that plague Fried's equipoise? Presumably, the idea is that this principle has a more robust epistemic threshold. Recall that on the more fragile view, a clinician would be obligated to provide A over B if she had a mere hunch that A was superior to B. On the present view, presumably the clinician could permit a patient to be randomized to A or B, even if she had a hunch that A was better, as long as that hunch is not supported by evidence that "ought to be convincing to colleagues." Presumably, if there is evidence that ought to be convincing to colleagues that A is superior to B, then it would be impermissible to allow that patient to be randomized to A. In fact, it may be medically irresponsible to allow randomization in that case. In effect, Miller and Weijer want to hold that clinicians who favor one intervention over another (e.g., A over B) can still allow their patients to participate in a study in which they will be randomized to A or B as long as doing so does not represent a medically irresponsible action, where "medically irresponsible" is a higher threshold than the standard of providing what the individual clinician actually believes is optimal care.

5.7.3 A Dilemma for the Clinical Judgment Principle

Miller and Weijer's proposal reflects the profound influence of the idea that the central issue to be resolved in research with human participants is to reconcile the duties of the individual clinician with respect for the welfare of the individual patient. Appealing to a more robust epistemic threshold is supposed to allow them to have the cake of locating uncertainty in the mind of the individual clinician even after having eaten the cake of avoiding the problems that plague the conception of equipoise that refuses to die.

But their view seems to face a difficult dilemma. Recall that Miller and Weijer find clinical equipoise deficient because it does not adequately address the duty of care of physician-researchers who must exercise their judgment and discretion in order to advance the agent-relative welfare interests of their patients to the best of their ability. But what is the relationship of the clinical judgment principle to the expert's duty of personal care? Either the clinical judgment principle is weaker than the clinician's morally and legally recognized professional duty to her individual patient or it is not. If it is weaker, then Miller and Weijer's own view can be rejected for not addressing what they regard as the central problem to be resolved, namely, reconciling research participation with the clinician's actual duty of personal care. If it is not weaker, then it is unclear how their position on this question differs from Fried's and therefore avoids the deep problems that his view faces (§5.6).

Although this dilemma can be easily stated, it cannot be easily addressed. We can amplify these concerns by revisiting the extent to which locating the focus of moral uncertainty back in the head of the medical expert recapitulates one of the very problems that clinical equipoise was developed to resolve—failing to recognize disagreement among experts as a kind of uncertainty that clinical trials ought to address (§5.6.3).

5.7.4 Conflicts over What Is Medically Irresponsible

A second major problem with this approach helps to flesh out the concern raised in the previous section. In particular, because the clinical judgment principle locates the relevant uncertainty in the head of the individual clinician, it cannot cope with situations in which expert disagreement runs so deep and is so polarized that the various sides question whether the care recommended by the others is ethically responsible.

Imagine a case in which some clinicians not only favor one intervention (e.g., A over B) but regard the other as medically irresponsible. Imagine further that other clinicians favor a different intervention (B over A) and regard the other as medically irresponsible. Now imagine further that each individual physician bases their judgment, not on a mere hunch, but on medical evidence that each regard as of sufficient credibility that it ought to be convincing to their colleagues. On the view articulated by Miller and Weijer, a protocol that would randomize individuals to these interventions could be approved by an IRB because such a body would correctly judge that clinical equipoise obtains—there is honest disagreement in the expert medical community about the relative merits of these interventions.

However, on Miller and Weijer's view, no clinician could permit her patients to enroll in such a study because doing so would violate the clinical judgment principle. That is, proponents of A would argue that it is medically irresponsible to allow their patients to be randomized to B and proponents of B would argue the same about being randomized to A. By reintroducing uncertainty in the mind of the individual clinician, Miller and Weijer's view recapitulates the same problems that we saw in §5.6—it prohibits socially valuable research without making anyone better off in the process.

5.7.5 Epistemic Humility

Miller and Weijer might argue that these last two objections misunderstand the force of the "ought" in the second condition of their principle of clinical judgment. In this view, if at least a reasonable minority of expert clinicians regard the evidence in support of A as sufficiently compelling that it ought to convince their colleagues, and a different group of at least a reasonable minority of experts believes the same about B, then both groups *ought* to update their beliefs and adopt the view that both treatments are above the threshold of medically responsible care. In other words, responsible medical professionals should show a modicum of epistemic humility in the face of such disagreements. Although this is a promising response, it suffers from several problems.

First, and most importantly, urging epistemic humility does not vindicate the importance of embracing uncertainty in the mind of the individual physician; it makes it irrelevant. This is because once we have established that clinical equipoise exists, we have established that there is sufficient evidence

to support A and sufficient evidence to support B that reasonable experts "ought" to regard randomization to each of these interventions as being consistent with competent or morally acceptable medical care. But, in this case, all of the real moral and epistemological work is being done by clinical equipoise and by an auxiliary claim that when clinical equipoise exists, reasonable clinicians ought to, in some sense, recognize the validity of the expert judgments of their honest and informed colleagues.

Second, this auxiliary claim is itself a substantive position that may seem plausible, but as a descriptive claim it need not be true and as a normative claim it requires substantive defense. In other words, it is not clear that it is irrational or unethical for different individual experts who are fully aware of all of the relevant medical evidence to draw conflicting treatment recommendations from that same set of evidence. The reasons for this claim take us beyond the scope of the current argument, and I will return to this issue briefly in the next chapter. But all that matters for our present purposes is that if clinical equipoise obtains and that is sufficient for the auxiliary claim, then Miller and Weijer's position adds nothing that was not already present in Freedman's view. On the other hand, if clinical equipoise is not sufficient for the auxiliary claim and if experts do not adhere to it in a particular case—if they regard the opposing view as medically irresponsible—then Miller and Weijer's view faces the objection we explored in §5.6; it would prohibit the conduct of a study that has significant social value without advancing anyone's interests in doing so.

5.7.6 Clinical Equipoise and the Particularities of Individual Patients

A second response might be to say that the analysis I have provided so far misconstrues the role of clinical equipoise and fails to take seriously the respect in which Miller and Weijer regard it as inadequate. In particular, Miller and Weijer argue that clinical equipoise only addresses the agent-neutral interests of participants; it does not and cannot address the agent-relative interests of study participants. So, this reply runs, clinical equipoise must be augmented by the judgment of a clinician who has a duty of care toward the *individual* patient in question.

The problem with this reply is that it misconstrues the role that clinical equipoise can and ought to play in research ethics—clinical equipoise

need not be limited to the agent-neutral interests of participants. Miller and Weijer correctly note that a trial protocol must be written at a certain level of generality, prior to an encounter with any particular patient, and that the question of whether or not a study would begin in and be designed to disturb clinical equipoise plays an important role in evaluating such protocols. They are also correct to note that inclusion and exclusion criteria are defined at the time the protocol is written and that it is important that these capture a realistic population of patients. Nevertheless, we can concede that they are also correct that if some patients could present with such a unique history—with characteristics that were not anticipated in the protocol's exclusion criteria—then the risks of study participation for that individual could be unreasonable. All of this is correct, as far as it goes.

But it is a mistake to think that just because IRBs must use clinical equipoise to determine whether or not to approve a particular protocol prior to the enrollment of individuals, that is the only place that clinical equipoise can be applied. To apply clinical equipoise at the level of individual patients, we need only ask whether, for each individual from whom consent is sought, experts who favor one intervention for patients with this condition would also regard that intervention as superior to the other alternatives *for this particular patient*. In other words, would those experts who favor treatment A over B for patients with this condition also prefer A over B for *this particular* patient? Similarly, would experts who favor B over A prefer B over A for *this particular* patient? If so, then it is permissible to randomize that patient to either A or B. Notice that it would be morally permissible even if each of these experts regards the evidence in favor of their preferred option as so strong that providing anything else violates Miller and Weijer's principle of clinical judgment.

Miller and Weijer appear to assume that questions about the unique medical history of particular individuals would have to be answered by a single individual and that that individual is the individual researcher. But this assumption is unnecessary. For example, imagine that after receiving IRB approval, a study begins to recruit participants. Each participant is evaluated by an expert who favors intervention A over B to determine whether in fact A would be an appropriate intervention for this person. This expert would determine whether, given the unique medical history of the person before them, there is any reason to think that A would pose unreasonable risks to this person (i.e., whether being given A is inconsistent with this person's agent-relative interests). Each participant is also evaluated by a second expert

who favors intervention B (and so on if there are additional interventions). If one or more of these experts finds that a particular person should not receive the intervention that they tend to favor for patients of this type, then such a person could be excluded from the study, or could be prevented from being randomized to that intervention if there are others (e.g., B and C) that are regarded as not unduly risky for this individual by the experts who regard each of those interventions as best for patients of this type. A design of exactly this type (only each expert is replaced by a computer model of the considerations that they regard as relevant to their clinical assessments) is described in Kadane (1996).

In such cases, no expert is asked to alter her beliefs in light of the conflicting judgments of other experts. Each is asked to make a medical judgment that best advances the interests of the patient before them. Nevertheless, no single individual expert need be uncertain about the relative merits of the interventions in question and, in fact, each can regard the views of the others as representing irresponsible medical care. This demonstrates how, contrary to the claim of Miller and Weijer, clinical equipoise can be used to regulate both the approval of the study protocol and the inclusion of individual participants and that clinical equipoise is sufficient to safeguard the agent-relative interests of individual patients.

In summary, then, Miller and Weijer's view is least objectionable when it is interpreted in a way that simply uses the existence of clinical equipoise to determine what individual clinicians ought to believe. To the extent that their view deviates from the requirements of clinical equipoise it recapitulates some of the problems that plague Fried's view. Ultimately, the analysis presented here shows that their argument for departing from Freedman's position rests on an unreasonably narrow understanding of how conflicting professional judgments can be used to evaluate both study protocols and the inclusion of individual study participants. The framework that I defend in the following chapter illustrates how a principle similar to clinical equipoise can address the concerns that motivate Miller and Weijer's departure from clinical equipoise without recapitulating the errors of Fried's view.

The arguments of this and the previous several sections provide strong reasons to reject any view in which uncertainty in the mind of the individual clinician is treated as a necessary condition for ethically acceptable research. Because this view is often treated as the only way to fill out the content of the template for the appeal to uncertainty, it is often assumed that the failure of this position demonstrates that there is a moral dilemma at the heart of

medical research. In truth, it shows only that this is a misguided way to fill out the details of the template.

In this section I argued that Freedman was aware of some of the limitations of treating the relevant uncertainty as located in the mind of the individual clinician. His claim that the relevant uncertainty should instead be treated as a function of beliefs of different experts in the medical community is important, and the view that I develop in the next chapter incorporates this insight. As we will see, however, the view I defend goes farther and rejects the commitment that Freedman shares with these other positions, namely, that the normative ground for the appeal to uncertainty is to reconcile role-related obligations of medical professionals with the demands of clinical research.

5.8 Purely Research-Related Risks

5.8.1 No Uncertainty about Purely Research-Related Risks

Even if it is possible to fill in the content of the template outlined in §5.4 in a way that avoids the problems discussed so far, it might be argued that this establishes that research can be organized to avoid a moral dilemma only if we limit ourselves to the interventions to which participants will be allocated. The next objection holds that there is, nevertheless, a dilemma at the heart of all research in which study participants are exposed to risks that derive from procedures or interventions that are necessary to advance the scientific aims of research and which are not offset by the prospect of individual benefit to participants. In other words, sometimes medical research requires tests or procedures that are performed solely to advance the purposes of research. They are necessary because they play a role in generating the data a study requires to assess the chosen endpoints or because they contribute to some other purely research-related desiderata, such as controlling bias. The worry, therefore, is that research that exposes participants to risks that are not offset by the prospect of direct medical benefit to those same participants poses a moral dilemma because it requires those participants to sacrifice their own welfare for the greater good.

This argument can be formulated in two ways. In §5.8.2 I present the version that focuses on the moral obligations of clinicians and researchers. In §5.8.4 I present a more general version that focuses on what it is rational for

potential study participants to choose. The latter version of this argument is of special interest since it is widely seen as grounding the claim that research participation represents an instance of the prisoner's dilemma.

The goal of this section and the next is to demonstrate that these arguments presuppose a conception of individual interests that is unjustifiably narrow. This view of individual best interests produces a conception of the clinician-researcher's duty of care or fiduciary duty that is so restrictive that it would rule out as inappropriate activities that are widely regarded as ethically permissible in the very area from which it is supposedly derived, namely, clinical medicine. As a result, I demonstrate in §5.9.1 this conception of the researcher's moral obligation to study participants is unjustifiably paternalistic. A parallel argument in §5.9.3 holds for the conception of individual interest presupposed in the claim that research participation constitutes a version of the prisoner's dilemma.

Together these arguments reveal the extent to which the first dogma of research ethics contributes to a conceptual ecosystem in which the appearance that research participation requires tragic choices is almost inescapable. Because Freedman's conception of clinical equipoise accepts this dogma of research ethics—it frames the point of the appeal to uncertainty as reconciling the clinician's duty of personal care with the demands of sound science—the problems discussed in this and the following section reveal important shortcomings in Freedman's conception of clinical equipoise. Together, these arguments illustrate the importance of finding an alternative normative ground for the appeal to uncertainty and reconsidering the first dogma of research ethics.

5.8.2 The Clinician-Centered Formulation

The *clinician-centered formulation of the argument from purely research-related risks* begins with the claim that a great deal of medical research involves practices, procedures, or interventions that are "not clinically indicated" (Wertheimer 2010, 9). These are interventions that would not be performed on a person in the context of direct medical care. Rather, they are provided because of the contribution they make to some important aspect of a research study. For example, in order to measure concentrations of a drug in a participant's blood, a study protocol may require study-related blood samples at regular intervals. In order to measure the effect of a drug on a

tumor, the protocol may require multiple biopsies at pre-specified intervals. In more extreme cases, in order to ensure that study participants cannot tell whether they received the active intervention in a trial or the control intervention, some study participants may be exposed to sham procedures. In the most benign cases, these procedures may involve mostly theater—surgeons reading a script, making superficial incisions in a participant's skin, and pretending to insert an arthroscope into the participant's knee, for example. But in other cases, the sham procedure can involve drilling a hole in a participant's skull and inserting a cannula which will deliver the investigational drug to those in the active arm and a placebo substance to those randomized to the control group (London 2006b, see also London and Kadane 2002, 2003).

The second claim is that the provision of such procedures cannot be justified by any view that requires research participation to be consistent with the individual clinician's duty of personal care. In particular, if that duty is understood along traditional, Hippocratic lines, then the clinician cannot recommend any course of care that is less optimal than some other possible course of care. But purely research-related procedures are provided solely to advance the scientific goals of a trial and not to advance the interests of the individual participant. As a result, clinicians who act on their fiduciary duty to put the interests of their patients above all other concerns, and who act on their duty to provide optimal courses of care to each patient, will not be able to support participation in any trial that exposes participants to such purely research-related risks.

To put matters in terms that link it more directly to the template outlined in §5.4, the risks and burdens of purely study-related procedures cannot be justified by the presence of uncertainty *no matter where it is located*. The risks and burdens of study-related procedures are usually not subject to the relevant kind of uncertainty—clinicians are not likely to be agnostic about whether such procedures align with and advance the interests of study participants. Rather, the opposite is likely to be the case—their risks and burdens are known and not reasonably seen as being offset by the prospect of direct benefit to the patient. Similarly, no clinician is likely to hold that the option of being exposed to such procedures in the course of a clinical trial has the same expected value for a patient's health interests as the option of foregoing study participation and receiving medical care directly. If everyone agrees that a given intervention or procedure carries risks and burdens that are not offset by the prospect of individual benefit, then there is also no

clinical conflict about their relative therapeutic, diagnostic, or prophylactic merits.

Even if uncertainty can bridge the divide between the clinician's duty of personal care and the demands of scientific research when it comes to the provision of alternative medical treatments or investigational interventions that are being tested as candidates for treatment, the objection currently under consideration holds that it can't play this role for purely study-related procedures. If a study protocol requires a set of blood draws or biopsies that would not be required in the context of normal medical care, then it is unlikely that even a reasonable minority of expert clinicians would regard those procedures as potentially beneficial to the individual trial participant. But if such uncertainty does not obtain, then we cannot appeal to the existence of uncertainty to reconcile study participation with the fiduciary duties of caregivers. Therefore these procedures appear to pose a dilemma for research ethics.

5.8.3 Compromising the Duty of Personal Care

The objection from study-related risks relies on a contingent feature of research, since different studies involve purely study-related procedures or interventions to varying degrees. In principle as well as in practice it is possible to design valid studies in which the relative merits of a set of interventions are explored without exposing participants to purely study-related procedures or interventions. This would be the case, for example, if the merits of these interventions are compared only on the basis of endpoints and measures that are routinely used in the course of delivering those interventions in clinical practice. Nevertheless, most studies with human participants do expose participants to purely study-related procedures that carry some risk or degree of burden. When this is the case, such research would be regarded as ethically impermissible on any view that requires its conduct to be consistent with the clinician's duty of personal care, understood as optimizing the medical interests of individual patients.

Even advocates for clinical equipoise seem to accept the conclusion of the argument in §5.8.2. In particular, the proponents of what is called "component analysis" restrict the scope of the equipoise requirement to interventions that are provided with "therapeutic warrant" (Weijer 1999, 2000; Weijer and Miller 2004). This includes interventions whose diagnostic, prophylactic, or

therapeutic merits are in question and under scrutiny in a particular study. As such, all purely research-related aspects of a study must be assessed on terms that reflect the weighing of different interests, as reflected in the definition of reasonable risk outlined in §5.2.4.

But adopting two standards for assessing research risks, as component analysis does, is a tacit admission that it is not possible to reconcile all aspects of clinical research with the clinician's duty of personal care. Since exposing individuals to procedures that, as Wertheimer puts it, are "not clinically indicated" cannot be reconciled with the clinician's duty of personal care, then critics can insist that component analysis shows that it is not possible for research to proceed on terms that are consistent with the clinician's duty of personal care. If what we "ought" to do is reasonably limited to what we "can" do (if "ought" implies "can"), then the proponents of component analysis must admit that it is permissible to carry out research on terms that diverge from the clinician's duty of personal care.

I have explicitly formulated the argument of this section as applying to the project of reconciling research participation with the role-related obligations of caregivers. This is an important objection and, as I argue in §5.9–5.10, it reveals a genuine problem for views that accept the first dogma of research ethics. Before turning to that discussion, however, it is important to present an alternative formulation of this argument that seems to have even broader scope and even more important implications. In particular, if it is true that research in which individuals are exposed to purely research-related risks cannot be reconciled with the clinician's duty of personal care, and if the clinician is seen as the fiduciary of the interests of the individual patient, then it seems to follow that participation in any such research is against the interests of individual participants and so not a rational choice for those individuals.

5.8.4 The Participant-Centered Formulation

In the previous three sections I explicated what I called the clinician-centered formulation of the argument from purely research-related risk*s*. In this section I introduce a related version of this argument that I refer to as the *participant-centered formulation of the argument from purely research-related risks*. What makes this formulation appear to be distinctive is that it seems to bypass an appeal to the role-related obligations of health professionals altogether, holding instead that research participation is fundamentally

inconsistent with the individual participant's concern for her own welfare. In other words, for any individual who is primarily concerned with her own medical best interests, clinical research appears not to be a rational choice.

Something like this argument seems to motivate the assertion of Menikoff and Richards that "tragic choices [are] involved in designing a system for research on human subjects" (2006, 19). Tragic choices are required because:

> Doing research involves *intentionally* exposing persons to risks, and *not for the primary purpose of treating them or making them better*, but rather to answer a research question. And, given the sorts of things that are commonly done in research studies, being a research subject in many cases will indeed be a bad choice for someone who is mainly concerned about his or her own best interests. (18)

If a person is "mainly concerned about her own best interests," then she will avoid participating in research because such participation so frequently involves being exposed to interventions, practices, or procedures that expose participants to burdens and risks without the offsetting prospect of direct, personal benefit.

5.8.5 Is Research Participation a Prisoner's Dilemma?

The idea that research participation is antithetical to the best interests of participants entails that if those individuals are choosing rationally, they will do all that they can to avoid research participation. At a social level, this creates a kind of paradox: although we all want to benefit from advances in the standard of care brought about by the conduct of well-designed research with human participants, none of us wants to be such a participant. The postulated moral conflict at the heart of medical research thus manifests at the social level in the form of a serious social dilemma.

The claim that medical research poses a social dilemma has been made by several scholars. David Heyd (1996) argues that research participation poses a social dilemma that is "reminiscent of the Prisoner's Dilemma" (193) because each potential participant would prefer to receive care directly from their clinician than to participate in a randomized clinical trial. If each person pursues what is in their individual interest, it forecloses advances in medical understanding. But, in order to agree to participate in a

randomized clinical trial, an individual would have to choose an option that is not as good as the available alternative from the standpoint of her narrow self-interest.

Alan Wertheimer makes a similar argument:

> Hence, we face a form of prisoners' dilemma. Suppose that the best available information suggests that it is 60% likely that intervention X is superior to intervention Y. Although it is in the *ex ante* interest of each individual not to participate in research and to simply receive X, it is in the interest of many others (including future persons) that a sufficient number participate in research to determine whether X is superior to Y with a greater level of certainty. Moreover, even if it were 50/50 as to whether X is superior to Y, it would be a bad choice to enter such a trial if one has to undergo procedures that were not clinically indicated or one were otherwise inconvenienced by participation. (2010, 9)

Wertheimer's claim that research with human subjects has the basic structure of a prisoner's dilemma draws on two sets of considerations that we have examined so far. The first (discussed in §5.6) is the idea that equipoise, conceived of as uncertainty in the mind of the individual expert, is fragile and evanescent—it will rarely obtain and even when it does it will not persist until the conclusion of a trial. The second is the idea that clinical research often involves tests or procedures that are not aimed at the medical best interests of participants. In both cases, Wertheimer argues that it is against the interest of potential participants to participate in research.

Wertheimer's formulation of the claim that research ethics requires tragic choices reveals the close connection between the clinician-centered and the participant-centered formulations of the argument from purely research-related risks. In particular, both arguments rely on a particular conception of the relationship between rational choice, welfare, and the health interests of the individual. In the participant-centered formulation, rational choice is equated with choosing the option that best advances the interests of the agent, where those are equated with that person's narrow health interests. The same view is presupposed in the clinician-centered formulation to the extent that it relies on the traditional, Hippocratic conception of the caregiver's duty of personal care. On that view, the duty of personal care requires clinicians to choose the option available to her that best optimizes the individual's medical best interests.

Purely study-related procedures and interventions appear to create a dilemma for research ethics because they are inconsistent with the medical best interests of participants. After all, these procedures or interventions are used, not because of the prospect that they will help the individual participant, but because of the way they contribute to a scientifically sound or socially valuable study design. A clinician who is obligated to choose only interventions or procedures that advance the narrow medical best interests of the patient before her cannot choose to expose individuals to such interventions. Similarly, if we assume that rational choice requires individuals to choose options that are in their own medical best interests, then no individual would rationally choose to participate in a study in which she is exposed to burdens and risks that are not offset by the prospect of direct, individual benefit.[7]

In the following section I deal with each of these arguments in turn. In §5.9.1 I argue that the clinician-centered argument from purely research-related risks is unjustifiably paternalistic and that we have independent grounds to reject this conception of the relationship between the duty of personal care, health, and patient welfare. I then argue in §5.9.2 that we have equally strong grounds to reject the more general position that individual rationality somehow requires individuals to choose only acts that optimize their narrow medical best interests.

[7] Although the *clinician-centered* and the *participant-centered* formulations of the argument from purely research-related risks are closely connected, the nature of that connection might differ, depending on how one approaches a larger set of questions. For instance, what we might call the *strongly role-related argument* holds that clinicians are obligated to advance the narrow medical interests of individuals for whom they are responsible because of the special role-related duties of caregivers. For example, one might argue caregivers have a special obligation to focus on patient health interests because of the centrality of health to their social role.

The *weakly role-related argument* holds only that caregivers are obligated to advance the best interests of individuals as those individuals understand them. Here, the clinician's focus on the narrow, medical best interests of individuals does not derive from anything internal to their professional role. Instead, it derives from (a) the deeper claim that in order for individuals to make rational choices they must choose the option that best advances their interests and from (b) the further claim that in matters of health, this necessarily involves choosing the act that optimizes their narrow health interests. If individuals understand their best interests as extensionally equivalent to whatever is in their narrow health-related interests, caregivers would inherit this focus on the patient's narrow medical interests.

As I proceed here, my critique of the clinician-centered argument dispenses with the strongly role-related version of this position (§5.9.1) and my critique of the claim that research participation is not in the narrow self-interest of participants dispenses with the weakly role-related argument (§5.9.2 and 5.9.3).

5.9 Well-Being and the Life Plan of Persons

5.9.1 Arbitrarily Restricting Individual Liberty

In §5.8.2 we saw that the clinician-centered argument from purely research-related risks is predicated on a fairly traditional, Hippocratic understanding of the clinician's duty of personal care. As a form of patient-centered consequentialism, it holds that "Physicians should promote the medical best interests of patients by offering optimal medical care; and the risks of prescribed treatments are justified by the potential therapeutic benefits to patients" (Miller and Brody 2002, 4). I now argue that we have independent grounds for rejecting this interpretation of the physician's duty of personal care. Rejecting this understanding of the duty of personal care removes one formulation of the argument which holds that there is a dilemma at the heart of research ethics.

Ironically, the grounds for rejecting the traditional, Hippocratic interpretation of the duty of personal care stem from applying it to the realm from which it is supposed to be derived, namely, clinical medicine. In fact, it is surprising that the Hippocratic conception of the duty of personal care persists in research ethics since the rejection of this view was one of the main drivers of contemporary medical ethics.

For the duties of Hippocratic patient-centered consequentialism to be aligned with patient interests it must be the case that health and health-related interests are the highest and most authoritative of the patient's interests (Goldman 1980). Although this is often the case, it is not always—and so not necessarily—the case. Patients sometimes have interests that take priority over their strict medical interests or that so color and shape those interests that it is difficult to disentangle their strict, medical interests from the larger set of interests that define their particular life plan. As a result, the larger contours of a person's distinctive life plan can lead them to make decisions that are at odds with what the Hippocratic clinician believes to be in their medical best interests.

To illustrate this point I want to focus, for the purposes of the present argument, on particular aspects of clinical medicine that bear structural similarities to purely research-related interventions or procedures. In particular, there are a range of practices in which patients undergo risks and burdens in clinical medicine solely for the purpose of assisting other people. Some examples involve relatively minor burdens and risks, such as blood donation.

Others involve more significant burdens, such as bone marrow donation. Still others involve even more significant burdens and medical risks, as when clinicians use their medical knowledge and skill to remove an organ or a portion of an orgen, such as a kidney or a lobe of the liver, from one person and transplant it into another.

If saving the lives of people in medical distress is an important project in a person's life plan, then the narrow medical or health risks of blood donation, organ donation, and other such procedures must be evaluated in light of the contribution that these activities make to the welfare of that same agent. Because some of these very acts are performed in both medical and research contexts, we can make the following direct argument. If the risks and burdens associated with drawing a person's blood violate the clinician's duty of personal care when performed in the context of a clinical trial, then those same risks and burdens must violate the duty of care when performed in the context of donations to be used by others in need. By modus tollens, because it is not impermissible for patients to donate blood for the purpose of advancing the interests of other people in the clinical context, it is not impermissible for study participants to donate blood in a clinical trial for the purpose of generating valuable information that is required to advance the interests of other people.

This argument demonstrates that the Hippocratic interpretation of the duty of personal care is more restrictive than the way that very same duty is interpreted in clinical medicine. Moreover, this is the same duty in both contexts. So, if the interpretation of that duty that is used in the clinician-centered argument is correct, then it would also rule out bone marrow donation, living organ donation, medical quarantine, and routine vaccination since all of these medical procedures impose some burdens or risks on one person for the purpose of generating a benefit that accrues to others. Since these activities are not regarded as inconsistent with the clinician's duty of care, then the we must reject the formulation of the clinician's fiduciary duties that animates the clinician-centered argument.

I have focused on cases that have a structural similarity to clinical research, because the Hippocratic conception of the duty of personal care requires medical procedures to advance the narrow medical interests of patients. Procedures such as vasectomy and tubal ligation are often not performed to rectify a medical pathology—to heal or alleviate pain or suffering or to restore what Daniels (1985) calls typical species functionality. Rather, those procedures are performed in order to assist individuals

in carrying out life plans in which they wish to engage in sexual activities without having to worry about procreating in the process. Those procedures would be regarded as ethically impermissible under the Hippocratic conception of the duty of personal care because they expose patients to medical risks and burdens to achieve goals or purposes that derive, not from addressing physical ills or medical pathologies, but from the goals of the individual's larger life plan.

Similarly, cosmetic procedures expose patients to risks that are not necessarily in the strict medical best interests of patients. Reshaping the contours of a fully functional nose, cheek, chin, breast, belly, and so on, are optional undertakings, often driven entirely by aesthetics. Many may question the wisdom of undergoing such procedures, and others may hold that because they are ethically optional undertakings there is no duty to use scarce resources to pay for them. But these are not the issues in question. On the argument we are considering here, offering such procedures to patients would be unethical because doing so violates the physician's duty of personal care—such procedures are intended, not to restore functioning or to treat disease, but to achieve aesthetic ideals.

The moral permissibility of medical procedures performed on patients for the benefit of others, or to advance goals other than a patient's strict medical interests follows from the rejection of medical paternalism. This was, in part, a rejection of the idea that the medical profession's specialized knowledge of health and disease was sufficient to understand the way that health or its absence influences patient welfare (Goldman 1980, 156–230). If health and the avoidance or amelioration of disease are sovereign values, the highest goal for any rational patient, then clinicians would have special insight into patient welfare in virtue of their special medical knowledge. But if health and the avoidance or amelioration of disease cannot necessarily be presupposed to be a person's highest goal or sovereign value, then which medical care best advances—or is most likely to frustrate—the interests of patients must be determined for each patient in light of that person's larger life plan.

The rejection of medical paternalism involved the recognition that the value of a state of affairs or of an outcome for a patient is not solely a function of that person's narrow medical interests; it depends on how those states or outcomes are situated relative to a patient's larger life plans (Goldman 1980). The very idea that a patient could have the right to refuse unwanted medical care—to withdraw a ventilator even when it is certain to fulfill its proper medical function of sustaining and extending that person's life—requires

recognition that health states that one person may regard as valuable and worth experiencing might be regarded by others as undignified and worth avoiding.

The reasonable diversity of life plans entails that although some individuals would not want to take any degree of personal risk or bear any burden to advance the health interests of others, other people view this as a calling and an avenue through which to express important values such as love, compassion, charity, solidarity, or reciprocity. Although some people would forego the prospect of extending their life if the means of doing so were painful, protracted, invasive, or risky, others often decide differently. Similarly, there are differences of judgment about the wisdom and value of reshaping one's body for motives other than the restoration of prior form or typical functioning and whether the attending risks are reasonable in light of the expected benefits, if any. This reasonable diversity in judgment about the risks and burdens of common medical practices stems from the reasonable diversity of life plans. This point is a concrete illustration of the respect in which the narrow technical perspective of medicine is incomplete (§5.3). How risks or burdens to a person's narrow medical interests impact that individual's welfare or wellbeing depends on their relationship to the projects and plans in that individual's larger life plan (see also §5.9.2 and §5.11). This information derives, not from the technical expertise of medicine, but from the reflective self-understanding of the individual whose interest medicine is expected to serve.

The rejection of medical paternalism was not a rejection of the idea that physicians and other health professionals have a fiduciary duty to individual patients. That duty is morally sound and important. Rather, the rejection of medical paternalism was a rejection of the idea that the traditional Hippocratic interpretation of the clinician's duty of personal care is a morally appropriate model of the relationship between patient health and patient welfare. That view has been repudiated in clinical medicine—the very domain in which it is supposed to be sovereign—because treating health as a proxy for individual well-being misconstrues the nature of human welfare. Equating welfare with a person's health elides the texture and complexity of the diverse life plans individuals in a free society can reasonably embrace. But it also gives too much authority to the social role of the clinician. In both cases, it arbitrarily restricts the autonomy of patients.

In the grip of the first dogma of research ethics, the field has retained the traditional, Hippocratic conception of the duty of personal care even after

that view was repudiated in the context of clinical medicine. The upshot of the argument of the current section is that the claim that the Hippocratic conception of the duty of personal care accurately reflects the content of the clinician's fiduciary duty to the individual patient is false.

If the clinician's fiduciary duty is interpreted, instead, as a duty to advance an individual's medical best interests as those interests emerge within that person's life plan, then this duty is not necessarily inconsistent with the performance of purely study-related procedures—even if those procedures carry affirmative risks and burdens. Rather, the permissibility of these procedures will depend on the extent to which the individual in question regards them as necessary and proportional burdens undertaken in the course of advancing an important project or plan. This point provides a kernel of insight on which we will draw in our response to the deeper and more philosophical problem about the relationship between rational choice, individual interests, and health, to which we now turn.

5.9.2 Personal Risks Are Not Irrational

It is important that when Menikoff and Richards assert that research participation is often a bad choice they scrupulously state that it is a "bad choice for someone who is mainly concerned about his or her own best interests" (2006, 18). This addition might seem trivial, since it might seem to be trivially true that every individual is mainly concerned about his or her own interests. But, in the sense in which this statement is trivially true, it is not necessarily inconsistent with an individual's best interests to participate in research in which they are exposed to burdens and risks that are not offset by the prospect of direct medical benefit. The reason is simple: if a project is sufficiently important to an agent that advancing it is a personal priority, then undertaking risks that are necessary to further that project is consistent with advancing their best interests.

On the other hand, if we understand concern for one's interests in such a way that it excludes accepting affirmative risks to one's health or welfare in the course of activities that primarily benefit others, then this claim is not only *not trivial*, it is so strong that many life choices would also pose a moral dilemma—including the decision of a young student to pursue a career as a physician or a medical researcher! We can elaborate the points in these last two paragraphs one at a time.

For the sake of the argument, let's assume that there is a tight motivational and rational connection between what an agent has a reason to do, what an agent is motivated to do, and what is in an agent's best interests. In particular, let us grant that if *x* is in an agent's interest, then that agent has a reason to support or engage in *x* and, conversely, that if *x* is not in an agent's interest then that agent has a reason to discourage or avoid *x*. Given these assumptions, it is critical to clarify what it means for something to be in an agent's interests.

Consider first the idea that an agent has an interest in *x*—that *x* is in the interest of an agent—if *x* is a constituent of that agent's life plan or if *x* is an instrumental means of advancing a project or element of such a life plan. This way of conceiving an agent's interests dovetails nicely with our previous claims about a close connection between *x* being in an agent's interest, that agent having a reason to do *x*, and being motivated to do *x*. In fact, this connection seems almost trivially true since it basically says that agents have a reason and a motivation to support or pursue whatever is a constituent of, or an instrumental means of effectuating, their particular life plan.

On this view, a wide array of things can feature into the life plan of an agent: careers, hobbies, ambitions, social connections or affinities, personal relationships and affections. In each of these cases, what constitutes advancing the agent's interest need not directly involve or appeal to any aspect of that person's health, physical status, or psychological state. Pursuing a career as an engineer, for example, can involve long hours dedicated to understanding the principles that organize some domain of the physical world and developing the knowledge and the means to use that knowledge to build structures, synthesize materials, or design and construct some other form of physical system. Success in the pursuit of such a career involves achieving the excellences that are associated with understanding the relevant systems, creativity in design, implementation or construction, efficient and safe use of resources, and so on. In such cases, an individual's life plan can revolve around an activity—such as designing and constructing a large and complex structure—to such a degree, and can involve exposure to such a range of associated risks, that pursuing that person's goals and ambitions can come into conflict with that individual's narrow health interests.

Civil engineers often work on construction sites in which there are positive risks of injury or death. Chemical engineers handle chemicals that can cause blindness, injury, or death. Similarly, physicians and medical researchers risk contracting illness from their patients, whether through direct exposure or from accidents such as needle sticks.

Part of the refutation of medical paternalism involved precisely this insight—that, in many areas of life, the life plans of individuals elevate the pursuit of other goals or ends over the maximization of individual health. The point of this insight is not to deny that health is an individual interest; it is simply to dethrone the idea that health is an agent's sovereign interest, trumping all others. Once we recognize that a person's life plan can elevate accomplishing some goal or set of goals above advancing their narrow health interests, then we can no longer assume that expertise in medicine provides sufficient insight into a person's interests to warrant empowering clinicians to subvert the freedom and choice of individuals in order to advance those individual's narrow health interests.

On this view, although it is almost trivially true to say individuals have reason to act in their own interest, and to avoid acting in ways that are not in their interest, the fact that participation in a study might expose a person to risks that are not offset by the prospect of direct medical benefit is not sufficient to establish that study participation is against that person's interests. If it were, it would establish that being an engineer or a clinician or a researcher is also against a person's interests. But such claims are false because a person's interests are not defined by their direct physical or mental status, but by the larger contours of their individual life plan. To know whether accepting such personal risks is consistent with or conflicts with a person's interests we have to know how those risks relate to the projects and goals that define their individual life plan.

If helping others plays an important role in one's life plan, and if donating blood is a means of helping other people in need, then when such a person donates blood, they are advancing their interests (the goal of helping others after a natural disaster, for example) despite the fact that the blood draw exposes them to both risks and burdens. Likewise, if finding a cure for a disease is one of a person's goals, and if extra blood draws are necessary to run a scientifically sound study, then undergoing those blood draws as a participant in a study can be in a participant's interests.

For many people, activities in which they take on risks and burdens to themselves in order to help others is a normal feature of everyday life. For example, people in many faith communities are called to engage in community service activities. Volunteers repair homes, provide care to the sick, and perform other tasks that are attended by personal risks and burdens. Similarly, many people identify deeply with their professions, including medical researchers. But medical researchers are often in contact with

needles, blood, and pathogens that they can and sometimes do contract. These hazards are often not discussed in public discourse and so researchers and participants are treated as though their respective pursuits are structurally very different—researchers advance their own interests and those of the larger community by conducting activities in which participants are exposed to risks that are inconsistent with their medical best interests.

The point I am making here is that if we focus on the medical best interests of these individuals, both being a researcher and being a study participant can involve risks that are inconsistent with that individual's narrow medical best interests. In both cases, respect for individual welfare requires that those risks be minimized, and gratuitous risks should be eliminated altogether. The key point, however, is simply that the existence of risks to a person's strict medical best interests is not necessarily inconsistent with a person voluntarily accepting those risks in order to advance the plan that imbues their life with personal meaning and social significance.

When Heyd (1996), Menikoff and Richards (2006), or Wertheimer (2010) assert that it is not in an individual's interests to participate in a study that poses some affirmative risks or that requires enduring some burdens or inconveniences, they are asserting a claim that entails that it is also not in an individual's interest to take on the career of a medical researcher. Such a result, however, is absurd.[8] Many people are drawn to a career as a researcher precisely because they see it as a way to use a diverse mix of scientific, mathematical, and social abilities to advance a worthwhile individual and social

[8] This narrow position seems more palatable when it is paired with what looks like an innocuous ancillary assumption. This assumption is that being a researcher is not in a person's individual interests until it is attached to a significant salary or elevated to a particular social status. On this view, being a researcher on its own is not in an individual's strict interests, but being a researcher as a way of securing significant wealth or social status renders it consistent with that individual's strict personal interests.

But this ancillary assumption is far from innocuous. In particular, when individuals value wealth or social status then attaching those things to an undertaking represents a way to encourage people to value that undertaking. But this move saves the narrow conception of individual interest by appealing to the broader conception of what it is to be in an individual's interest to which this narrow theory is supposed to be an alternative. In particular, it isn't clear why doing *x* for the money is supposed to be easier to grasp as a rationale for engaging in *x* than doing *x* as an outlet for one's various talents and abilities, or because it contributes to a cause to which one is committed. Making money does not make a direct contribution to one's narrow health interests. Instead, it is either valuable as an end that one embraces for itself, or as a means to advancing the other ends that one embraces, including advancing one's health interests. But developing one's talents and abilities and pursuing one's larger life projects might make an affirmative contribution to one's physical and mental health. Even if it doesn't, developing one's talents and abilities is either an end in itself or a means of advancing other ends that one embraces. As a result, doing *x* because it advances a life project or represents the expression of one's talents and abilities seems no worse, and possibly better, as an explanation for why *x* is in one's interests than doing *x* for the money.

project. Many scientists identify so closely with the ends that they pursue that they have been willing to put their life on the line, whether as a subject in their own study (Altman 1972; Neuringer 1981) or as a researcher in dangerous contexts in which they could contract life-threatening disease or be subject to violence (Green 2014). These behaviors are not only permitted, they are often valorized.

In a community in which different individuals pursue a diversity of life plans, it is likely that many people are willing to accept affirmative risks to their own health if those risks will contribute to the knowledge needed to understand and ultimately alleviate suffering or disability associated with sickness, injury, or disease. In such cases, personal risks may be unwanted and not assumed lightly or without adequate safeguard, but insofar as they cannot be avoided and are tied to activities that are constitutive of or instrumental to a person's life plan, they are not necessarily inconsistent with that individual's best interests. Researchers and study participants differ in many morally relevant respects. In particular, participants are likely to face particular risks with greater certainty because they are part of an explicit and formalized research protocol. Nevertheless, researchers and participants can be symmetrically situated in their acceptance of risks to their narrow, medical interests, in order to advance meritorious social ends.

5.9.3 Study Participation Is Not a Prisoner's Dilemma

The argument in the previous section allows us to demonstrate precisely why research participation does not give rise to a prisoner's dilemma. To make this case, it is helpful to carefully lay out the structure of this particular social dilemma. Doing so reveals an interesting fact—that contrary to the assertions of those who make this claim, research participation is not in fact this type of strategic dilemma. Even so, once we distinguish the impact of participation on an individual's health interests from its impact on their overall interests, we can demonstrate that study participation can be a rational move to make in this kind of strategic situation.

Figure 5.1 contains a simple diagram that illustrates the structure of a prisoner's dilemma.[9] In this example, each individual has to make a choice

[9] Although this example focuses on two individuals, this is merely for convenience. Two-person prisoner's dilemmas can be scaled up to n-person prisoner's dilemmas without altering the results.

Each individual's best choice is to "defect," no matter what the other does.		**Individual 2** (Welfare score is bottom right number in each box.)	
		Cooperate	Defect
Individual 1 (Welfare score is top left number in each box.)	Cooperate	5 5	1 6
	Defect	6 1	2 2

Figure 5.1 Simple representation of a two-person Prisoner's Dilemma. Individual 1's welfare score in each outcome is indicated by the top left number in each box. Individual 2's welfare score in each outcome is indicated by the bottom right number in each box.

between participating in some collaboration or defecting. The boxes represent outcomes that result from the respective choices of each player and the numbers inside each box represent the "payoffs," "utilities," or "welfare score" for each party in that state of affairs. For convenience I have chosen integers to represent welfare scores, but, once again, the structure of the dilemma does not depend on our being able to assign specific numbers to individual welfares. The same dilemma emerges as long as the relative orderings of outcomes depicted in the table are preserved.

In the example from which this problem derives its name, two prisoners are being interrogated by the police. If they both cooperate, and keep silent, they go to jail for only one year. In the matrix in Figure 5.1, this is the cooperate/cooperate square in which each player receives 5 welfare units. But if one player keeps silent and the other defects, blaming his partner for their nefarious activities, then the defecting partner goes free (a score of 6 units) and the silent party goes to jail for the maximum sentence (a score of only 1

The number of parties involved is thus less important than the structure of the problem they face—that is the central issue.

unit). If each defects and exposes the other they go to jail for slightly less than the maximum sentence (2 units each).

From the standpoint of the players in this game, the outcomes produced when both cooperate are preferable to those produced when both defect. The dilemma arises from the fact that in this situation, there is no third-party "social agent" who gets to choose which outcome actually happens. In other words, there is no dictator who can force these two individuals to cooperate. Instead, each party has to choose for themselves whether to cooperate or to defect and each is expected to choose in his or her own interests, where the numbers in the box represent the agent's interests. Given the assumption that each individual makes a rational choice when they choose the option that best advances their own interests, it is not rational for either individual to choose to cooperate. No matter what Individual 1 chooses, Individual 2 receives a higher welfare score from defecting and no matter what Individual 2 chooses, individual 1 receives a higher welfare score from defecting. This result does not depend on complicated solution concepts from game theory such as the Nash equilibrium. It hinges simply on the fact that, for each individual, defecting produces more individual welfare than cooperating regardless of what the other person chooses. In the formal language of game theory, the choice to defect "dominates" the choice to participate.

Prisoner's dilemmas can arise in a wide range of contexts. As such it is best not to think about the motivating story of the two prisoners and to focus instead on the relationships between the payouts in the matrix. Any social interaction in which the interests of the parties are accurately modeled by welfare payouts with these relationships will face this thorny problem. When commentators assert that research participation is a prisoner's dilemma they mean, not that the parties are prisoners of some kind, but that the choices they face have consequences whose values are accurately modeled by the numbers in Figure 5.1.

Despite its allure, the claim that research participation is a prisoner's dilemma is false. When someone refuses to participate in a study they continue to experience or undergo whatever course of care or state of affairs is the status quo. As a result, the coordination game facing potential participants in medical research is represented in Figure 5.2. For simplicity, the status quo for each party is represented as "0," as neither a gain nor a loss. As a result, the party who defects does not stand to gain something by defecting, at least when measured relative to the status quo. She may gain something relative to

	Participate	Don't participate
Participate	x+y x+y	x 0
Don't Participate	0 x	0 0

Figure 5.2 The coordination game reflecting the strategic decision of parties who must decide whether or not to participate in medical research. Here, x represents the direct positive or negative impact on the health of the agent from participating in the study and y represents the agent's personal valuation of the information that the study is designed to produce.

the other player, however, depending on what is involved in participating in research.

In light of our discussion in the previous section we can represent the effects of participation with two variables. The direct impact (positive or negative) on the health of the agent expected from participation is represented by x. The valuation of participation, as represented by the agent's valuation of the information the study is expected to produce, is represented by y. In Figure 5.2, y is only present in the cell in which both players participate. This is intended to mark the idea—reflected in the original claim that research participation has the form of a prisoner's dilemma—that the social benefits of research require cooperation of other willing participants. In that respect, y only materializes if a sufficient number of individuals are willing to participate that the study can be run to completion with sufficient power. To maintain the simplicity of the representation, therefore, it is best to think of the parties in this example as small groups.

The nature of the strategic situation represented in Figure 5.2 depends on the values that x and y take, but none of these values produces a prisoner's dilemma. The situations associated with the different values of x and y is represented in Figure 5.3. If x offers sufficient prospect of direct personal benefit to participants ($x > 0$) then there is no dilemma and no coordination problem; everyone prefers to participate rather than not to participate. If study participation involves affirmative risks and burdens (such that $x < 0$), and no agent values the information the study is likely to produce to such a degree that it would compensate for those personal risks or burdens ($x + y < 0$) then there also is no dilemma; nobody participates because nobody thinks the study is worth the risk.

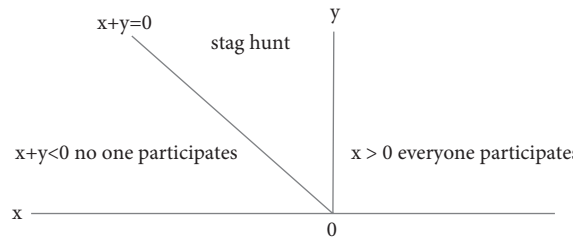

Figure 5.3 The strategic structure of the game represented in Figure 2 depending on the values of x (direct health benefits or burdens to individual participants from study participation) and y (the value to the agent of the information a study is likely to produce) relative to the status quo (represented as 0).

If the study involves affirmative risks to participants ($x < 0$) but the prospective participants value the information that the study is likely to produce enough that they are willing to accept those risks ($x + y > 0$), then the game has the form of what is called a stag hunt (Skyrms 2004). In a stag hunt, each individual prefers to cooperate only on the condition that others cooperate as well. Unlike the prisoner's dilemma, defecting is not the dominant course of conduct in a stag hunt. In that regard, the stag hunt is a coordination problem and not a social dilemma. If the agent is convinced that there are other agents who also value generating the evidence the study is designed to produce, and so will participate on condition that others participate as well, then joint cooperation is a rational choice. Put in slightly different terms, under the circumstances just described, participating in research in which the participant will be exposed to some burdens or risks that are not offset by the prospect of direct medical benefit remains an equilibrium of the game and therefore a rational choice for a rational agent.

In a pluralistic community, different agents may have different attitudes toward the same study. Some agents may regard the associated risks as unreasonable in light of the way participation fits into their larger life plan. In contrast, other agents may view study participation as contributing to a worthwhile project or reflecting important aspects of their personal life plan. In such cases, the society is faced with a coordination problem—if enough people embrace a life plan that is advanced by generating information that will help to understand, treat or prevent a debilitating disease and are willing to participate in research, then society need only provide assurance to such

individuals that if they participate then enough like-minded people will follow to generate socially valuable information.

The upshot of the argument in this section is that accepting affirmative risks to one's narrow health interests is a routine part of pursuing a distinctive life plan and that, as a result, accepting purely study-related risks when they are necessary to promote an activity that a participant values and wants to promote is not inconsistent with that agent choosing in a way that advances her best interests. From this I showed that research participation is not a prisoner's dilemma.

The arguments of the last two sections show that the participant-centered formulation of the argument from purely research-related risks does not reveal a dilemma at the heart of all research that involves purely research-related interventions or procedures. As a result, any conception of the researcher's duty of personal care or fiduciary duty to the individual that focuses narrowly on that individual's health interests is unduly restrictive and unjustifiably paternalistic.

5.10 Against the First Dogma of Research Ethics

5.10.1 Hippocratic Duty Has Clear Content but Is Unjustifiably Restrictive

Together, the arguments in this chapter provide powerful reasons to reject any view that seeks to constrain research activities by requiring that they be consistent with the individual clinician's duty of personal care. We can express the cumulative force of these arguments in the form of a dilemma. Call this the *dilemma of determinate duties*.

One horn of the dilemma holds that if the caregiver's duty of personal care is interpreted in traditional, Hippocratic terms, as a form of patient-centered consequentialism, then it has independent content that places clear constraints on the practice of research. The problem is that this clarity of content is purchased at an unacceptable price: the limits that it imposes on research are unjustifiably restrictive. This standard would deny individuals who identify with the goals of a socially valuable activity (whether as study participants or researchers) the ability to accept any degree of personal risks or burden in the furtherance of that activity. Not only is such a standard repudiated in clinical medicine—the very context from which it

supposedly derives—but its implementation on a social scale would unreasonably restrict the array of life plans individuals in a free society are capable of pursuing and deprive communities of the social benefits generated from the willingness of individuals to adopt life plans that include activities that promote the common good.

If it is consistent with the clinician's fiduciary duty to permit blood donation, bone marrow donation, living organ donation, or cosmetic surgery, then it should be consistent with this same duty to permit study participants to accept those burdens or risks for the purpose of facilitating socially meritorious research. This suggests that when study participants identify with the goals of a clinical trial, study-related blood draws, tumor biopsies, and other procedures can be consistent with the duties and norms of caring medical practice.

5.10.2 Duty of Care That Respects Autonomy Lacks Independent Content

The second horn of the dilemma holds that if we interpret the duty of personal care in a way that is more aligned with how it is understood in medical practice, then we strip from that duty the independent content necessary to set determinant constraints on research risk. In part, this stems from the fact that limits on research risks have to be set prior to the point at which individual study participants are approached with the offer to participate. But if the judgment of the reasonableness of risks requires knowledge of a person's larger life plan, then IRBs would lack the information they need to apply this standard in the evaluation of study protocols.

The problem also stems from the fact that the limits we impose on individual decision-making in clinical practice may be overly permissive. In particular, caring medical practice includes respect for the wishes of competent patients to refuse life-sustaining medical care, hastening the patient's death. It is not clear how this would translate into an analogous standard for limiting risk in research. Would it be permissible for study participants to knowingly hasten their own death in pursuit of research-related objectives? One worry is, thus, that this standard would not provide any substantive constraint on what could be offered to study participants. Instead, it would rely solely on the procedural constraint of whether participants are willing to consent to

whatever the study protocol demands, no matter the magnitude of the risks that might entail.[10]

Once it is recognized that reasonable people, with different life plans, can value various health states in different ways, and can have divergent preferences about the means of promoting those health states, then medicine cannot be treated as a form of productive knowledge that has a complete, self-contained understanding of individual patient interests. Rather, medicine is incomplete in the sense that it must look to the larger values and interests of the individual person to fill out the picture of how various health states, and the means to achieve them, fit into an individual's conception of the good life (§5.3.1–2). Making this move brings the content of the clinician's duty of personal care into better alignment with the way that duty is understood in clinical medicine. But it does this at the cost of stripping that duty of its operational content when it comes to regulating research.

The refutation of medical paternalism was not just a repudiation of arbitrary power that had been vested in the hands of physicians; it was a recognition that such power was arbitrary precisely because it did not track the larger interests of individuals as free persons who exist for themselves and not for the purposes of others. Although health is an important good, the place of that good in an individual's larger life plan is ultimately determined by the shape and contours of that larger conception of flourishing and what constitutes a good life. Medicine is an incomplete guide to welfare since the place of health in a person's larger life plan cannot be determined solely by the technical principles of medicine. Rather, that information has to be provided from outside medicine, from the practical judgment of the autonomous individual.

These problems illustrate a point I made in §5.3. Professions are relatively narrow bodies of knowledge that range over distinct domains. But how those domains impact the interests of individuals, and what justice demands of them in a free society, are issues that fall outside the narrow confines of professional obligations. Recognizing this point and repudiating the first dogma of research ethics are essential to understanding how to regulate research risks in a way that is consistent with the requirements of the egalitarian research imperative.

[10] For a defense of this position see Rajczi (2004). For a critique see Rid and Wendler (2010).

5.11 Against the Second Dogma of Research Ethics

5.11.1 Utilitarian Assumptions Are Not Necessary

The appearance of a deep conflict at the heart of research with humans is encouraged by the perception that purely research-related risks are necessarily antithetical to a person's interest and, therefore, that the only way to justify their presence in research is by ensuring that they are offset by the expected benefits to future beneficiaries of research. The idea that risks to some must necessarily be traded off against benefits to others is part of the second dogma of research ethics.

The second dogma of research ethics rests on an unnecessarily strong assumption about what is required in order to ensure that research has social value. In particular, even if a frankly utilitarian approach to research risks would be sufficient to promote socially valuable research, it does not follow that it is necessary to achieve this end. But it is the latter, stronger claim that is needed to show that the research enterprise poses a deep conflict with the rights and interests of study participants.

In this chapter I have argued that the perception that research participation is necessarily inconsistent with concern for individual welfare is mistaken. As we attenuate this perception of conflict, we also attenuate the idea that the only way to advance medical science is to be willing to sacrifice individual interests for the benefit of future persons. After all, if scientifically sound and socially valuable information can be generated without requiring compromise in any relevant value—if it does not require breaching participant rights or sacrificing participant welfare—then pursuing the requirement of social value would not be inconsistent with *any* ethical perspective. Even the strictest absolutism about rights and values could support the vision of research participation articulated in the egalitarian research imperative.

I have argued here that when agents adopt socially meritorious ends—ends that involve aiding or assisting others—and those ends cannot be advanced without the agent being exposed to personal risks and burdens, such exposure is not necessarily antithetical to the overall interests of those individuals. This is not to say that it is acceptable for such risks to be gratuitous—it is not. Nor is it to say that such risks should be lightly undertaken. It is to say that when a person donates bone marrow to save

the life of another, the pain, inconvenience, and risks of complications that are part of the donation process should be reduced and properly managed, but that they are not antithetical to the welfare of the person making the donation.

To bring what might seem like a counterintuitive claim into sharper relief, it is helpful to contrast two scenarios and the moral principles that might be used to assess the reasonableness of risks in each.

5.11.2 The Principle of Proportionality

Sam is a firefighter who takes great pride in having the strength, stamina, courage, experience, and knowledge necessary to fight fires. Fighting fires, saving property, and rescuing people from hazardous situations are a source of pride for Sam and an outlet through which Sam both experiences certain personal goods and makes a social contribution. Sam's position is paid for by the Township. The Township values Sam as one of its members. It also bears various kinds of responsibility in the case that Sam is injured, becomes disabled, or is killed in the line of duty. Some of these responsibilities are financial. But others relate to the way the Township values people—whether the Township values its workers and its members as free and equal persons or whether it treats them as disposable tools.

In some contexts, Sam appears to be willing to accept greater personal risks than the Township thinks reasonable. The Township therefore wants to limit the risks to which Sam is exposed when acting as an employee of the Township. It relies on the principle of proportionality as a guide to limiting those risks.

Principle of Proportionality: A condition for the acceptability of risks within an activity is that those risks must not be disproportionate in comparison to the goods they are necessary to generate.

Both the Township and Sam agree that this principle allows for different levels of risk depending on the nature of the activity Sam is undertaking. For example, the maximum permissible risk in the course of saving property is lower than the maximum permissible risk in the course of trying to save a person. Even in the latter case, however, there is a limit to what constitutes an acceptable risk to Sam.

5.11.3 The Principle of Utility

In contrast, Pat is an artist whose life plan revolves around exploring personal expression through various visual media. Pat happens to be physically fit and, because of a shortage of firefighters in Municipality, Pat is conscripted into service as a firefighter. Although Pat, like Sam, is physically capable of performing the relevant tasks, Pat, unlike Sam, has no interest in serving as a firefighter. Every moment is an exercise in drudgery and boredom interrupted by moments of sheer terror for Pat.

Municipality invokes the principle of utility to justify selecting Pat to serve as a firefighter.

Principle of Utility: It is permissible to perform an act that decreases the welfare of one person as long as doing so produces a sufficiently large increase in the welfare of others.

In particular, Pat's life plan is set back by service as a firefighter, because Pat values the excellences and experiences of making and appreciating art and every moment spent as a firefighter is a moment taken away from the activities in which Pat finds fulfillment. Nevertheless, Pat's ability to fill this role better than the other candidates is regarded as generating sufficient benefits to others to justify this reduction in Pat's welfare.

Although the principle of utility might entail both the social value requirement and the principle of proportionality, the social value requirement and the principle of proportionality do not necessarily entail the principle of utility. The reason is that the injunction to ensure that the risks a person undertakes in the course of a voluntary pursuit are reasonable in light of the goods they are seeking to produce does not entail the permission to sacrifice the welfare of any agent to promote the good of others.

As we saw in chapter 2, both proponents and critics of the idea that there is an imperative to carry out research conceived of study participation on the model of Pat—as someone whose interests are abrogated or sacrificed for the benefit of others. Implicitly, however, researchers have been treated on the model of Sam—someone who accepts personal risks in the course of an occupation with the noble goal of advancing medical knowledge in order to be a benefactor to humanity. On that model, as we noted previously, the risks that professionals such as clinicians and researchers incur to their health in the course of pursuing their personal and professional objectives are largely

ignored. The two dogmas of research ethics perpetuate this asymmetric framing by hard-coding the idea that research necessarily involves trade-offs of the kind to which Pat is subjected.

The egalitarian research imperative that I defend here is an imperative to create a system of knowledge production that gives a symmetric treatment of researchers and study participants on the model of Sam. The goal is to create a social system in which both researchers and participants can participate freely as an avenue through which they have the opportunity to generate an important public good. We have more closely approximated this ideal in the case of researchers and other medical health professionals. In chapter 7 I argue that part of the benefit of a system of prospective review of research before bodies of diverse representation is that it brings us closer to achieving this ideal for research participants as well. Reconceiving research ethics on the model defended in this book would move us much closer to this goal. The point is not to say that researchers and study participants are somehow exposed to the same level or kind of risks—for this is surely not true. Rather, the point is to create a system of voluntary participation in which no person is conscripted like Pat, forced to sacrifice their welfare for the greater good, and in which every participant can be treated on the model of Sam.

In a society in which people pursue diverse life plans, different individuals will be drawn to study participation for different reasons. Whether they are likely to contribute to such research will depend on our ability to provide credible public assurance that the endeavor in which they participate has significant social value, that their participation contributes to the production of an important public good, and that no other party to this cooperative scheme has the ability to co-opt it for their personal, parochial ends (see §5.9.3 and chapters 4 and 7). Although we have seen several arguments that seek to establish that research cannot be organized on such terms, none of those arguments is compelling.