

ABSTRACT

Justified Closure of Inquiry: A Non-Reductive Account

Kraig W. Martin, Ph.D

Mentor: Jonathan L. Kvanvig, Ph.D.

This dissertation is concerned with answering the following question: In virtue of what is one justified in closing inquiry? A closure of inquiry in an act whereby one intentionally stops some investigation and does so with the intention of not investigating in the future. I argue that on a maximization model of practical rationality, one is fully justified to close inquiry whenever one is morally obligated to close inquiry, or one is practically and morally justified in closing inquiry, or one is morally and evidentially justified in closing inquiry and one is justified in believing that closing inquiry is as good as any other means to the achievement of one's ends (perhaps because one is justifiably worried that any defeaters one finds will be misleading defeaters). On a model of practical rationality in which an act is practically rational for one only if one is justified in believing that the act is as good as any other means toward the achievement of one's ends, then one is fully justified in closing inquiry whenever one is morally obligated to close inquiry, or one is morally and practically justified in closing inquiry

Justified Closure of Inquiry: A Non-Reductive Account

by

Kraig W. Martin, B.S., M.A.

A Dissertation

Approved by the Department of Philosophy

Michael D. Beaty, Ph.D., Chairperson

Submitted to the Graduate Faculty of
Baylor University in Partial Fulfillment of the
Requirements for the Degree
of
Doctor of Philosophy

Approved by the Dissertation Committee

Jonathan L. Kvanvig, Ph.D., Chairperson

C. Stephen Evans, Ph.D.

Alexander R. Pruss, Ph.D.

Trent G. Dougherty, Ph.D.

Ryan S. King, Ph.D.

Accepted by the Graduate School
May 2014

J. Larry Lyon, Ph.D., Dean

Copyright © 2014 by Kraig W. Martin

All rights reserved

TABLE OF CONTENTS

Chapter One: When an Act is a Closure of Inquiry	1
Is Closure of Inquiry ever Justified?	3
The Necessity of Epistemic Normativity	8
Epistemic Justification and Closure of Inquiry	13
Fully Justified Closure of Inquiry	18
Chapter Two: The Necessity of Epistemic Normativity	22
Instrumental Rationality, Practical Rationality, Epistemic Rationality, and Morality	24
Every Act is Either Practically Rational or Practically Irrational and Morally Justified or Immoral	28
The First Fully Reductive Possibility: Reduction to Practical Rationality	34
Denying Moral Normativity	35
Morality as a Species of Instrumental Rationality, or Morality as a Constraint on the Permissible Ends of Instrumental Rationality	45
The Second Fully Reductive Possibility: Reduction to Morality	46
The Conjunction	49
The Disjunction	55
Epistemic Normativity and Justified Closure of Inquiry	56
Chapter Three: Closure of Inquiry and The Sufficiency of Knowledge	58
An Argument for The Sufficiency Principle	61
Millar on Reflective Knowledge	63
Potential Counterexamples to TSP	68
Jessica Brown and the Diligent Surgeon	72
Epistemic and Alethic Justification: A Helpful Digression	79
Conclusion	83
Chapter Four: Closure of Inquiry and the Sufficiency of Epistemic Justification	85
Premise Two	86
Premise Three	86
Premise Five: The Easier Objections	88
Premise Five: The Radical Externalist Worry	92
Premise Five: The Moderate Externalist Worry	96
Premise Seven	102
Premise Nine	105
The Upshot of EJ''	106

Chapter Five: Fully Justified Closure of Inquiry	110
A Non-Reductive Theory of Justified Closure of Inquiry	112
Evidential Justification without Practical Justification	116
Practical Rationality and Utility Maximization	121
Conventions, Etiquette, and Aesthetic Normativity	130
Conclusion	131

ACKNOWLEDGMENTS

This dissertation would not have been possible without the help of a good many people. To begin, I would like to thank my graduate student colleagues; I can't imagine a better group. Ross Parker, Nate Jackson, Pete Younger, Logan Gage, John Giannini entered the program in same year that I entered, and all of them helped me become a better philosopher. Many other of my graduate student colleagues, in classes both before me and after me, helped me become capable of writing this dissertation.

Along with my colleagues, an incredible group of mentors went beyond the requirements of the job in the time and effort they spent in my training. My committee—Ryan King, Trent Dougherty, Steven Evans, Alex Pruss, and Jonanthan Kvanvig—all spent significant time both preparing my prior to the writing of this dissertation and in the early and late drafts during the actual writing process. Of these, I want to especially thank my director, Jonathan Kvanvig, who spent many early morning with me discussing these ideas, exchanged many emails correcting my early thoughts, and read many drafts of many pages, some of which made this final version, and many of which fell to wayside. I am deeply grateful for his training, his insight, and his wisdom.

The philosophy department at Baylor is fantastic place, and I will always be thankful to God that I was fortunate enough to spend time with the people who make up that department. It is a place not only of professionalism, high standards,

incredibly smart philosophers, and a deep sense of community, but also a place of shared vision for doing Christian philosophy, a task of great importance.

Finally, I wish to thank my wife, Susan. I was thirty-two years old, having never taken a philosophy class, when the writings of Dallas Willard and Alvin Plantinga wooed me to the pursuit of philosophy. For Susan to enter this journey with me, a journey of near-poverty and no guarantees for the future, tells one all one might need to know about what an amazing person she is.

To my wife Susan, who gave so much that I might pursue philosophy

CHAPTER ONE

When an Act is a Closure of Inquiry

Poor John. He couldn't get any sleep. He went to bed, and immediately began wondering whether or not he remembered to turn his alarm off. He really wanted to sleep late in the morning, so it was important to him that he remembered to turn it off. After checking it the first time, he immediately started worrying that he looked at it wrong. So he checked again. He lay back down and thought, "If I could look at wrong once, I could do it twice. I can't really be sure that it is off." So he looked again. And a third time. And a fourth time. His wife was merely irritated at first, but soon became worried. John couldn't stop checking the alarm, over and over again. Eventually, he fell into an exhausted heap on the floor. After three days of this, John was involuntarily committed. Poor John. He just couldn't stop investigating.

It seems like there are circumstances in which we can (and perhaps even should) close down an investigation. In general, after we've found something, we can stop looking for it. After I find my shoes, I stop looking for them. In fact, many people might call me stupid, or even crazy, if I kept looking for my shoes even after I had found them. When we investigate, we are looking for something—truth, maybe, or knowledge, or reflective knowledge. It seems plausible to suppose that, when we find truth, or knowledge, or reflective knowledge (whatever it is we are looking for), we can—maybe even ought to—stop looking.

This dissertation is itself a series of investigations. My question is this: When are we justified in closing down an investigation? There is some number of propositions that are plausible answers to that question. Regarding each of those propositions, we are inquiring into whether or not they are the case. We are looking for true propositions that are answers to the question, and are general enough to be theoretical in structure.

When someone closes inquiry, as I use the term, someone acts intentionally. As I use the phrase, it is what one does when one decides to quit looking for evidence regarding the truth of some proposition. When one closes inquiry, one does not expect to reopen it (although one may). Scientists are no longer looking for evidence into the truth of *the earth is at the center of the universe*, because they know that it is false.¹ Inquiry into the matter is closed. When I was a kid, I wondered if *Corvettes are faster than Porsches*. I never did find out, but I quit caring, so I quit inquiring. I have closed inquiry into *God exists*; I am satisfied that he does. At one point, most scientists closed inquiry into *Newtonian physics accurately predicts movement in the universe*. They thought they knew that it was true. At some point, they reopened inquiry into that proposition.

I take it that inquiry can't be closed unless it is first opened. As such, one does not close inquiry into the truth of a proposition unless one first inquires into the truth of the proposition. Further, closing inquiry does not occur every time one stops looking for evidence into the truth of some proposition that one at one time looked for evidence regarding; dying, for example, is not the act of closing inquiry.

¹ Here, and throughout this document, I will use italics as a means of drawing attention to specific propositions.

Rather, closing inquiry, although one need not be conscious of it *as* closing inquiry, must be an intentional cessation of evidence gathering.

In this introductory chapter, I first consider two arguments that closure of inquiry is never justified.² If closure of inquiry is never justified, then there is no use looking for a theory about when it is justified. Responding to these arguments helps clarify the notion of closure of inquiry, and gives us hope that there is a theory of justified closure of inquiry to be found. After responding to these arguments, I provide a roadmap for the structure of the rest of the dissertation.

Ultimately, I argue in favor of a non-reductive theory of justified closure of inquiry. By non-reductive, I mean that (at least) morality, epistemic normativity, and practical rationality must be included as fundamental constituents of any successful theory of justified closure of inquiry.

Is Closure of Inquiry ever Justified?

Despite the fact that it seems that we often do legitimately stop looking for evidence regarding the truth of some propositions, there is a concern that closure of inquiry is never justified. I here assume that virtually all of what we know, perhaps literally all of what we know, is fallible knowledge; that is, there is some small chance that it is false. Suppose that we had a chance at infallible knowledge, and that the acquisition of this knowledge would be cost free. Shouldn't we always accept cost-free infallible knowledge? It is plausible to suppose that we should. If closing inquiry regarding a proposition requires us to refuse all offers of infallible

² I owe a special debt of thanks to Alexander Pruss, who presented a version of both of these arguments to me in a very helpful email exchange.

knowledge regarding that proposition, then it seems as though closure is never justified, since it would be a mistake to refuse to accept cost-free infallible knowledge. Call this the *Always-Accept-Cost-Free-Infallible-Knowledge* (AACFIK) argument.

AACFIK is instructive. First, it is not clear to me that we should always accept cost-free infallible knowledge. Suppose that you are standing before God. He says to you, "Ask me anything you want, I'll tell you, cost free." Suppose that you respond with, "I'd really like to know if you exist, God. If I can get two questions in, I'd also like to know if I exist." Even ignoring the opportunity cost (you might have asked something else), there seems to be something wrong with your response. I know that I exist. I don't need to ask God about that one. And it's rude to ask him if he exists. Maybe the lesson is not that it is sometimes okay to decline cost free infallible knowledge, but that cost free knowledge (infallible or otherwise) is metaphysically impossible.

Another instructive aspect of AACFIK is that it helps us see more clearly that a closure of inquiry need not entail that one never decide to reopen inquiry. A closure of inquiry is intentional cessation of evidence gathering, but it does not entail that one can never change one's intention. Suppose I want to make chocolate chip cookies, and I begin to inquire into whether or not I currently have vanilla in the pantry. (It helps to imagine a dirty pantry.) When I find a bottle of vanilla, I don't keep rummaging through the pantry, and I intend to no longer rummage. When I find that the bottle is empty, I begin to rummage again; that I can reopen inquiry does not entail that I never closed it. Since that one has closed inquiry does

not entail that one will not reopen inquiry, it is not true that one who has closed inquiry is committed to rejecting future offers of cost free infallible knowledge, assuming such offers are metaphysically possible.

That one can reopen inquiry highlights the fact that there is a difference between a cessation of inquiry and the mere suspension of inquiry. I noted earlier that not all cessations of inquiry were closures of inquiry, because death does not entail a closure of inquiry. A death is not an intentional cessation of evidence gathering. Another important aspect to the intentionality of a closure of inquiry is that a closure of inquiry entails not only the intention to stop looking for evidence but also the intention not to look for evidence in the future. One who intentionally closes inquiry but does so without the intention to not look for future evidence merely suspends inquiry. One intentionally ceases an inquiry into a proposition when one decides to quit looking for evidence regarding the truth of that proposition, and one intends not to reopen inquiry later—although one may well change one's mind; we often change our mind about many of our intentions.

A second argument against the possibility of legitimate closure of inquiry is also instructive. It goes like this: If I am justified in closing inquiry into a proposition, p , then I am justified in not looking for evidence into any proposition that includes p as a conjunct, unless the evidence I am attempting to gather has no evidential bearing on p . But there is some proposition that includes p as a conjunct in regards to which I am not justified in failing to gather evidence that bears evidentially on p ; so there must be no cases in which I am justified in closing inquiry.

(Call this argument *The Conjunction*). I will illustrate The Conjunction with an example.

Recall poor John, who was unable to stop investigating whether or not his alarm was turned off. In a modified case, suppose John is an alarm clock repairman, and he sleeps in a loft above his shop. He has a particularly good reason to need a full night's sleep tonight, and there are one hundred alarm clocks in his shop. Of each clock, he is nearly certain that it is turned off. (Perhaps he even knows regarding each clock that it is turned off). If there were only one clock, he would be justified in not checking it. In regards to the conjunction, however, he is not justified in closing inquiry; he ought to go check. How does he check that *A1&A2&A3...&A100 are turned off*? He goes to check to see if *A1 is turned off*, and if *A2 is turned off*, and so one. If John was really justified in closing inquiry concerning *A1 is turned off* then he would be justified in not looking for evidence regarding the relevant conjunction. But he is not justified in not looking for evidence regarding the relevant conjunction.

The Conjunction fails to show that closure of inquiry is never justified, but it is instructive regarding what a closure of inquiry is. When John is checking each clock, he is not looking for evidence regarding each conjunct; he is looking for evidence regarding the conjunction. Intentionality matters to a proper description of the evidence that one seeks. One can close inquiry into a proposition and still look for evidence regarding the truth of that proposition, as long as one is not looking for it *as* evidence for the truth of that proposition. We should say that one

closes inquiry into a proposition, *p*, when one intentionally stops (not merely suspends) looking for evidence (*qua* evidence) for or against *that* proposition.³

Suppose that when John picks up clock A1 and looks at the positioning of the alarm button, it is a psychological fact about him that he is not looking with the intention of finding evidence regarding *A1 is off*, but rather with the intention of finding evidence regarding *A1&A2&A3&....A100 are off*. In such a case, John is not investigating whether *A1 is off*. He is investigating the conjunction, not the conjunct.

Admittedly, there is still pressure to think that John is checking to see if A1 is off, and is intending to check to see if A1 is off. If you remain convinced that this is the case, then I offer the following distinction: some investigations are mere means to further theoretical goals, and some are not mere means to further theoretical goals.⁴ That one has closed (and not reopened) an inquiry into *p* entails that one will not look for evidence (*qua* evidence) regarding *p*, unless that evidence gathering is merely a means to a further theoretical goal. Since there is a principled difference between investigating *this alarm clock is turned off* and *all one hundred of these alarm clocks are turned off*, this distinction is not *ad hoc*. Since this distinction allows us to uphold the highly plausible prima facie claim that closing inquiry is sometimes justified (just think of poor John in the initial alarm clock case!), we should accept it.

³ The Conjunction is unsound either because its first premise is false (if “looking for” entails the intentionality necessary for an inquiry) or because its second premise is false (if “looking for” does not entail the necessary intentionality).

⁴ Again, thanks to Alexander Pruss for this distinction.

While the arguments against the possibility of legitimate closure of inquiry are instructive, they should not sway us from what seems obviously true. If one is sometimes justified in intentionally stopping an investigation, and doing so with the intention of not investigating in the future, then closures of inquiry are sometimes legitimate. And, the fact is, we often do justifiably stop investigating, and we sometimes do so with the intention of not investigating in the future. With this result, I launch an investigation: In virtue of what is closure of inquiry justified?

The Necessity of Epistemic Normativity

In chapter two, I argue that, whatever the correct theory of justified closure of inquiry, it must include a role for epistemic normativity. I have two arguments for this conclusion. The first argument is lengthy and ambitious. In this argument, I attempt to look at all of the domains of normativity that might plausibly be thought to explain when a closure of inquiry is justified. The most likely candidates, other than epistemic norms, are practical norms and moral norms. I show that no combination of the presence and/or absence of practical and/or moral justification can provide a plausible account of when it is legitimate to close an investigation.

In the process of making this first argument, I begin by explicating a conception of instrumental rationality in which both epistemic and practical rationality are a species of instrumental rationality. I make these distinctions for the sake of clarity, though nothing rides on them. I then consider all of the logical ways that practical rationality and morality might combine to provide the output for a single normative judgment considering the justificatory status of a closure of inquiry. I begin by assuming that a closure of inquiry could be (a) practically

rational/practically irrational/neither practically rational nor practically irrational and (b) morally justified/morally unjustified/neither morally justified nor morally unjustified. I immediately argue that every act is either practically rational or practically irrational, and every act is either morally justified or immoral. Thus, the logical space available for a theory of justified closure of inquiry that relies on practical rationality and morality is quickly and drastically reduced.

In defending the claim that every act is practically rational or practically irrational, I argue in favor of two theses: First, that if one is justified in believing of some act that it is as good as any available alternative in achieving one's relevant end, then the act is practically rational. Second, that if one is not justified in believing of some act that it is as good as any available alternative in achieving one's relevant end, then the act is practically irrational. Since one is always either justified or not justified in believing a given proposition, then every act is either practically rational or practically irrational. I also defend the claim that every act is either morally justified or immoral. I take this claim to be more intuitive, since every act is either morally permitted or not morally permitted.

With this result in hand, I begin to look at every possible combination of practical rationality/practical irrationality and moral justification/immorality/moral obligation that might explain when a closure of inquiry is justified. Some of these possibilities are fully reductive in that they attempt to explain justified closure of inquiry through a principle that makes use of only a single domain of normativity. I first consider a theory that reduces justified closure of inquiry to practical rationality. One who holds that one is justified in closing inquiry whenever it is

practically rational for one to do so would need to hold that no true propositions about one's being practically rational in closing inquiry can clash with true propositions about what one morally ought to do. Such a person might well make use of either an argument inspired by either expressivism or error theory.

Such a position entails either that statements of the form 'S ought to φ ' are always false (at least when the ought in question is a moral ought) or that such statements fail to express a proposition. With the first disjunct in mind, I consider the claim made by error theorists that there are no moral properties. Even supposing that there are no moral properties, it does not follow that all moral claims are false. I provide a view defended by Judith Jarvis Thompson as an example of a position one could hold while consistently denying the existence of moral properties and maintaining that there are statement of the form 'S ought to φ ' which are true, even when the ought in question is a moral ought.⁵ With the second disjunct, I discuss the problem of embedding; moral claims are often embedded in sentences or utterances that clearly express propositions. I then discuss Gibbard's response to the problem of embedding, and conclude that it is not sufficient to save the conclusion that moral claims always fail to express a proposition.⁶ In doing so, I borrow heavily from Ralph Wedgwood.⁷ Having dismissed the most plausible reasons one might have for thinking that no true propositions about one's being practically rational in closing inquiry can clash with true propositions about what

⁵ Thomson, *Normativity*.

⁶ Gibbard, *Thinking How to Live*.

⁷ Wedgwood, *The Nature of Normativity*.

one morally ought to do, I turn to the second possibility for a fully reductive theory of justified closure of inquiry.

MJ S is justified in closing inquiry concerning whether or not p iff S is morally justified in closing inquiry into whether or not p .

By modifying an argument formulated by Trent Dougherty in defense of the claim that each instance of epistemic irresponsibility/irrationality is just an instance of non-epistemic irresponsibility/irrationality, I formulate a strong *prima facie* argument in favor of the reduction necessary for MJ.⁸ If practical rationality can be reduced to a more fundamental normative concept without losing the ability to explain relevant data, then the reduction ought to be accepted.

1. If practical rationality can be reduced to a more fundamental normative concept without losing the ability to explain relevant data, then the reduction ought to be accepted.
2. Morality is more fundamental than practical rationality.
3. Practical rationality can be reduced to morality without losing the ability to explain the relevant data.
4. The reduction ought to be accepted

All that is necessary to show that this argument is unsound is a single counterexample to (3). I give such a case.

After considering and rejecting the fully reductive theories, I turn my attention to the truth functional ways that practical rationality can be combined with moral justification. I first consider the conjunction of practical rationality and moral justification.

PR \wedge MJ: S is justified in closing inquiry concerning whether or not p iff it is practically rational for S to close inquiry into whether or not p and S is morally justified in closing inquiry into whether or not p .

⁸ Dougherty, "Reducing Responsibility."

I argue that $PR \wedge MJ$ fails because one might be morally obligated to close inquiry on some occasion when doing so is not practically rational. In response, I consider

$(PR \wedge MJ) \vee MO$: S is justified in closing inquiry concerning whether or not p iff
a) it is practically rational for S to close inquiry into whether or not p and S is morally justified in closing inquiry into whether or not p , or b) S is morally obligated to close inquiry.

This principle makes no obvious use of the epistemic domain, and expresses a very plausible account of when justified closure of inquiry is justified. Nonetheless, I argue that either practical rationality is itself fundamentally dependent on the epistemic domain, or there are counterexamples to the theory. I then briefly consider and reject a theory that is constructed with the disjunct of PR and MJ.

$PR \vee MJ$ S is justified in closing inquiry concerning whether or not p iff at least one of the following is true: a) it is practically rational for S to close inquiry into whether or not p or b) S is morally justified in closing inquiry.

$PR \vee MJ$ fails for the same reasons that the left-to-right entailment of MJ fails; there are cases in which one is not morally forbidden from closing inquiry, yet one is not all-things-considered justified in closing inquiry.

Finally, I briefly present a second argument, one used by Fantl and McGrath, in favor of my conclusion that no satisfactory theory of justified closure of inquiry can exclude a role for epistemic normativity.⁹ If justified closure of inquiry were all about practical rationality and morality, it should not be possible to hold the non-epistemic domains steady, adjust only the epistemic domain, and get a different correct verdict on whether or not the closure was justified. But we are able to do just that.

⁹ Fantl and McGrath, *Knowledge in an Uncertain World*.

Epistemic Justification and Closure of Inquiry

Chapters three and four together comprise an argument in favor of:

EJ'' S's evidence justifies S in closing inquiry concerning whether or not p
iff S's belief that p (or S's belief that $\sim p$) is epistemically justified.

The following, which I call my Main Argument, is my argument in favor of EJ'':

1. If S knows p , then S is evidentially justified in closing inquiry concerning whether or not p . [premise]
2. Either S's knowledge that p or a fundamental constituent of S's knowledge that p explains why S is evidentially justified in closing inquiry concerning whether or not p . [premise]
3. S's knowledge that p can't explain why S is evidentially justified in closing inquiry concerning whether or not p . [premise]
4. A fundamental constituent of S's knowledge that p explains why S is evidentially justified in closing inquiry concerning whether or not p . [from (2) and (3)]
5. The fundamental constituents of S's knowledge that p are the following: p is true, S has epistemically justified belief that p , and S's belief that p is undefeated. [premise]
6. The feature of S's knowledge of p that explains why S is evidentially justified in closing inquiry concerning p is one of the following: p is true, S has epistemically justified belief that p , and S's belief that p is undefeated. [from (4) and (5)]
7. That p is true and that S's belief that p is undefeated cannot explain why S is evidentially justified in closing inquiry. [premise]
8. Therefore, that S is epistemically justified in believing that p explains why S is evidentially justified in closing inquiry concerning whether or not p . [from (6) and (7)]
9. If that S is epistemically justified in believing that p explains why S is evidentially justified in closing inquiry concerning whether or not p , then EJ'' is true. [premise]
10. Therefore, EJ'' is true. [from (8) and (9)]

In chapter three, I am only concerned with defending (1), which I call The Sufficiency Premise (TSP), and with showing the initial plausibility of EJ''.

In trying to show the initial plausibility of EJ'', I first discuss what I mean by evidential justification. Sometimes, one is justified in closing inquiry not because one has come to know that about which one had been inquiring, but because one quit caring about whether or not the proposition is true or false. In such cases, it

makes sense to say (loosely), that, if one is justified in closing inquiry, it is not one's evidence that justifies the closure. In other cases in which one is justified in closing inquiry, one is justified not because (or, perhaps, not merely because) the answer no longer matters, but because one has found the answer. In the second case, but not the first, one is evidentially justified in closing inquiry. Later, in chapter five, I discuss the relationship between evidential justification, practical rationality, and morality.

After providing a brief argument in favor of TSP and showing its initial plausibility, I consider potential counterexamples to TSP. First, I consider and reject Alan Millar's claim that reflective knowledge, not merely knowledge, is necessary for a legitimate closure of inquiry.¹⁰ Millar argues that knowing p is insufficient for reaping the benefits of knowing p , because one needs to know that they know p before that can enjoy the inquiry related benefits. I show that his requirement would make all evidentially based closure of inquiry unjustified. It is far too strong.

I then consider the similarities between TSP and a principle that Jeremy Fantl and Matthew McGrath have defended.

KJ If you know that p , then p is warranted enough to justify you in φ -ing, for any φ .

KJ, while it entails TSP, is stronger than TSP. I consider two potential counterexamples to KJ and argue that they are not problems for TSP. The first counterexample I consider, Barron Reed's jelly bean case, is instructive because it shows how KJ is stronger than TSP; while it is a potential counterexample to KJ, it is

¹⁰ Millar, "Why Knowledge Matters," 64.

not a problem for TSP.¹¹ The second potential counterexample, Jessica Brown's surgeon case, is the right kind of case to be a potential counterexample to both KJ and TSP.¹² I give reasons for thinking that KJ can survive the counterexample, but don't commit myself to the conclusion that it can. I argue that TSP, on the other hand, clearly has the resources to handle the case.

After defending TSP, I make a brief digression in order to provide myself with the resources to handle an important worry about EJ". In some cases, it seems extremely plausible to think both that one is justified in believing *p* and that one is not justified in closing inquiry into whether or not *p*. Following Jonathan Kvanvig, I argue for a kind of justification that, while it resides in the epistemic domain, is not epistemic justification; I also follow Kvanvig in calling such justification alethic justification.¹³ Merely that one is justified in believing *p* does not demonstrate that one has epistemic justification for *p*.

In chapter four, I finish the work of arguing for EJ" by defending (3), (5), (7), and (9) of The Main Argument. After completing that defense, I provide one additional upshot of EJ"—it does much of the theoretical work of KJ without falling prey to a particularly compelling counterexample.

A large section of chapter three is spent defending (5), the claim that the fundamental constituents of *S*'s knowledge that *p* are *p*'s truth, *S*'s epistemically

¹¹ Reed, "A Defense of Stable Invariantism."

¹² Brown, "Subject-Sensitive Invariantism and the Knowledge Norm for Practical Reasoning."

¹³ Kvanvig, "Coherentism and Justified InConsistent Beliefs: A Solution," 30–31.

justified belief that p , and that S 's belief that p is not externally defeated. The most worrisome objection to (5) is the claim that epistemic justification is not needed for knowledge; instead, what is needed for knowledge is truth, belief, and some proper external relationship, such as reliably formed true belief, or true belief that is the result of proper functioning faculties operating in the environment for which they were intended. These views are externalist theories of knowledge.

I follow Fantl and McGrath in distinguishing between radical externalism and moderate externalism.¹⁴ I rely on my defense of TSP from chapter two, and argue that radical externalism is incompatible with TSP. This is because the truth of radical externalism entails the possibility that one could know p and be justified in believing not- p . Given that such an agent knows p , and given TSP, it follows that the agent would be evidentially justified in closing inquiry concerning whether or not p . But it is not plausible in such a case to suppose that the agent's evidence is such that the agent is evidentially justified in closing inquiry, at least not in the right way; the relation between what such an agent knows and what her evidence supports is backwards.

The truth of moderate externalism does not entail that one can know p and be justified in believing not- p , because the moderate externalist grants that justification is necessary for knowledge, but insists that whether or not one is justified in believing should itself be understood as obtaining when the appropriate external relations obtain. Nonetheless, moderate externalism, too, is incompatible

¹⁴ Fantl and McGrath, *Knowledge in an Uncertain World*. See chapter three.

with the truth of TSP because their reply to the new evil demon problem, when applied to cases of closure of inquiry, is implausible.

Given that externalism is incompatible with TSP, and given the plausibility of TSP, and my defense of TSP in chapter 2, we should dismiss externalism. Once we dismiss externalism, (5) becomes extremely plausible.

My defense of (7), like my defense of (5), largely mirrors an argument in favor of KJ that is defended by Fantl and McGrath. When we imagine the presence of truth, without the presence of epistemic justification, it seems clear that closure of inquiry is not evidentially justified. When we imagine the presence of whatever constitutes the Gettier condition without the presence of epistemic justification, it seems clear that closure of inquiry is not evidentially justified. When we imagine the presence of both of these conditions, without the presence of epistemic justification, it seems clear that closure of inquiry is not evidentially justified. When we imagine the presence of epistemic justification, it makes no difference as to whether truth and/or the Gettier condition is present—closure of inquiry is evidentially justified.

After a defense of (9), I provide an upshot of EJ"—it does much of the theoretical work of KJ without falling prey to a particularly compelling counterexample. Specifically, KJ entails that if S knows *p*, then S is justified in taking a bet that he is metaphysically certain that *p*. But there are many things we know about which we are not metaphysically certain, and we ought not bet that we are.

Fully Justified Closure of Inquiry

With a defense of the Main Argument, I have defended the claim one is evidentially justified in closing inquiry concerning p if and only if one is epistemically justified in believing p . In chapter two, I showed that, whatever justified closure of inquiry is, there needs to be a role for the epistemic domain. In chapter five, I argue that either the correct theory of justified closure of inquiry is non-reductive, or it is reduced by an elimination of the practical domain. In calling a theory non-reductive, I mean that (at least) morality, practical rationality, and epistemic rationality have a fundamental role in the theory.

The added insight that chapters three and four provide into what constitutes evidential justification to close inquiry allow us, in chapter five, to more cleanly distinguish between other kinds of justifications that fall short of being full, or all-things-considered, justification. When one is *fully justified* to close inquiry, the full normative picture—all normative considerations—justifies one in doing so. One can be evidentially justified in closing inquiry without being fully justified in closing inquiry because one might be morally obligated to continue inquiry. Whether or not one is fully justified in closing inquiry is sensitive to (at least) whether or not one is morally justified in closing inquiry, whether or not one is practically justified in closing inquiry (i.e., whether it is practically rational for one to close inquiry), whether or not one is evidentially justified in closing inquiry, and whether or not one is morally obligated to close inquiry. (Since one is never all-things-considered justified in doing anything that one is morally obligated not to do, moral normativity can trump practical and epistemic normativity).

With the necessary distinctions in hand, I look at twelve ways these normative domains might relate. I argue that one has full justification to close inquiry whenever one is morally obligated to close inquiry. In those cases in which one is morally permitted, but not morally obligated, to close inquiry, then one is fully justified in closing inquiry if the following is also true: 1) one is practically and evidentially justified to close inquiry, or 2) one is practically justified but not evidentially justified. In the latter case, then if one is justified in believing p (where p is the proposition into which one is, or was, inquiring) one merely has alethic justification to believe p . If one is morally permitted to close inquiry but lacks practical and evidential justification to do so, then one is not fully justified to close inquiry.

I then discuss the possibility of a scenario in which one is morally and evidentially justified in closing inquiry, but not practically justified in doing so. If such a scenario is not possible, then the correct theory of justified closure of inquiry is partially reductive; it can be stated without appealing to practical normativity. While I do not argue that is possible for one to be morally and evidentially justified in closing inquiry without being practically justified in doing so, I tend to think it is possible. Assuming it is possible, then either the correct account of practical rationality is dependent on the epistemic domain, or in some of those cases in which one is morally and evidentially (but not practically) justified in closing inquiry, one is fully justified in closing inquiry.

In making this argument, I attempt to show that a common assumption about practical rationality is false; that is, I attempt to show that it is false that an act is

practically rational for one if and only if the act maximizes one's expected utility. In short, I think such an account of practical rationality is too demanding. I also consider and reject satisficing views, such as those discussed by Slote and Hurka.¹⁵

The success of my argument, however, does not hinge on my ability to conclusively demonstrate the falsity of maximization models of practical rationality. Supposing that such models are true (which I doubt), then there are situations in which one lacks practical rationality to close inquiry, and yet is fully justified in closing inquiry. Specifically, there are cases in which closing an inquiry maximizes one's expected utility (or meets whatever demands the proper satisficing theory includes), but one is justified in believing that continuing inquiry will reveal only misleading defeaters regarding that proposition into which one is evidentially justified and morally justified in closing inquiry. An adequate account of fully justified closure of inquiry, in order to account for such scenarios, would need to include a role for the epistemic domain. In doing so, the theory would be non-reductive.

On a maximization model of practical rationality, one is fully justified to close inquiry whenever one is morally obligated to close inquiry, or one is practically and morally justified in closing inquiry, or one is morally and evidentially justified in closing inquiry and one is justified in believing that closing inquiry is as good as any other means to the achievement of one's ends (perhaps because one is justifiably worried that any defeaters one finds will be misleading defeaters). On a model of practical rationality in which an act is practically rational for one only if one is

¹⁵ Slote, *Beyond Optimizing: A Study of Rational Choice*; Hurka, "Two Kinds of Satisficing."

justified in believing that the act is as good as any other means toward the achievement of one's ends, then one is fully justified in closing inquiry whenever one is morally obligated to close inquiry, or one is morally and practically justified to close inquiry.

Finally, I consider the worry that there are domains of normativity that I have not considered. While morality, practical rationality, and epistemic rationality seem to be the three most influential types of normative constraints, they are not exhaustive. There are also, for example, societal conventions, etiquette, and aesthetic standards. It may be possible for these other normative standards to affect the status of a closure of inquiry in a way that a correct theory of justified closure of inquiry must account for.

In response to this worry, I explain that the non-reductive theory of closure of inquiry that I have defended should be viewed not as a fully developed theory of justified closure of inquiry, but merely as a "toy theory" or a heuristic. Rules of thumb are often very helpful in both practical and theoretical pursuits, and the non-reductive theory of justified closure of inquiry that I explicate should be viewed in that light

CHAPTER TWO

The Necessity of Epistemic Normativity

I begin my search for a theory of justified closure of inquiry by first considering the possibility that what justifies a closure of inquiry can be reduced to purely non-epistemic normativity. In considering this possibility, I will help myself to the use of some fairly common philosophical assumptions. Whether the truth of these assumptions is necessary for a defense of my developed theory will need to be addressed in due course. For now, I am going to simply take these assumptions as a starting point from which to investigate what it is that justifies one in no longer looking for evidence regarding some proposition.

I assume that there are exactly three distinct normative domains of the sort that influence the justificatory status of a closure of inquiry – the practical, the moral, and the epistemic.¹ The practical is governed by practical rationality, the moral is governed by morality, and the epistemic is governed by epistemic rationality. If these three domains of normativity exhaust the types of normativity which might plausibly be thought to regulate justified closure of inquiry, then the presence or absence of one of these justifications (or some combination(s) of the

¹ There are very likely more domains of normativity than just these three. Rules of etiquette and the constraints of the aesthetic domain, for example, might not fit into any of these three categories. Nonetheless, I propose to start by considering only the moral, epistemic, and practical domains. I do so because they are the domains which most obviously affect the appropriateness of ceasing an investigation. In my final chapter, I return to the possibility that other domains might enter the calculus.

presence or absence of these justifications) must explain when closure is or is not justified.

If one also assumes the controversial claim that the epistemic domain only governs belief (and not actions), then one might be inclined to think that the correct theory of justified closure of inquiry (as an action) need only appeal to practical rationality and/or morality. In this chapter, I am interested in debunking that very thesis. In other words, I will argue that closure of inquiry is not regulated solely by non-epistemic normativity.

I am going to present two arguments for this conclusion. In my first argument, I am going to look at all of the ways that practical and moral justification might combine. If a correct theory of justified closure of inquiry can be articulated using only these two domains of normativity, then it should be possible to express this theory (or some principle that is extensionally equivalent to this theory) by merely stating the nature of the relationship between practical rationality and morality that is necessary and sufficient for a justified closure. For example, I might be practically justified in closing inquiry but lack moral justification for closing inquiry. Alternatively, I might be morally justified and lack practical justification. After explaining the ways in which these two domains might combine, I'm going to look at all plausible principles that make use of these combinations to explain the presence or absence of justified closure of inquiry. I will argue that all of these principles fail.

Finally, I will argue that we should not be surprised that they fail. After all, if practical rationality and morality alone could explain the presence or absence of

justified closure of inquiry, then we would not be able to hold those factors steady and yet get a different verdict on whether or not one is justified in closing inquiry. But, I will argue, we are able to do just that.

In order to make progress in this manner, I need to first make a case for the different ways that practical rationality and morality might relate to one another. In order to do that, I need to say more about what practical rationality and morality are.

Instrumental Rationality, Practical Rationality, Epistemic Rationality, and Morality

Instrumental rationality is that type of rationality that has to do with achieving some end, regardless of the acceptability, moral permissibility, or rationality of the end itself. Thomas Kelly writes, “By *instrumental* rationality, I mean the rationality which one displays in taking the means to one's ends. Thus, if I have the goal of asking the speaker a question, and I know that I will only be able to ask the speaker a question if I raise my hand, then (all else being equal) it is instrumentally rational for me to raise my hand.”²

I will talk as if epistemic rationality and practical rationality are each a species of instrumental rationality. Neither of claims is uncontroversial.³ While I

² Kelly, “Epistemic Rationality as Instrumental Rationality,” 612.

³ Raz, “The Myth of Instrumental Rationality”; Kelly, “Epistemic Rationality as Instrumental Rationality.” Raz objects to the claim that all practical rationality is a species of instrumental rationality, and Kelly objects to the claim that epistemic rationality is a species of instrumental rationality. Both concede that the claim to which they object is not uncommon among philosophers. While the view that I assume here is not uncommon, I hope that I don't need to defend it. Neither claim is central to any argument that I employ, and all of the language that I use might be reformulated to drop these two assumptions.

admit to being sympathetic to both claims, I am hopeful that I do not actually need either one of them. Neither of them do any work (at least, no obvious work) in my argument.⁴

Here is my conception of epistemic rationality: a belief is epistemically rational if it is rational to believe relative to the end of believing truly.⁵ Understood in this way, epistemic rationality is a species of instrumental rationality. In Kelly's example, raising one's hand is what one is rational in doing relative to the end of asking the speaker a question. Instrumental rationality is characterized by the right sort of means/end relationship. When the means in question is a believing, and the end in question is believing truly, then the species of instrumental rationality that emerges is epistemic rationality.

I use the term 'practical rationality' to specify that species of instrumental rationality in which the means to the end is a doing.⁶ This allows for a principled

⁴ I do eventually make a claim about practical rationality that does work in my first argument, but I will defend that claim, and my defense of it does not need it to be true that practical rationality and epistemic rationality are both a species of instrumental rationality.

⁵ This does not entail that a proposition must be believed in order to be epistemically rational for one, nor that when two people have identical evidence in regards to some proposition that they must agree regarding the truth or falsity of the proposition or (at least) one of them is epistemically irrational.

⁶ This articulation of practical rationality does not preclude the possibility of having practical reasons for belief, as long as some beliefs (call them believings) can be doings. The baseball player has practical reason to believe he is going to get a hit; such a belief increases his chances of getting a hit. However, to the extent that his believing is responsive to such reasons, it might plausibly be understood as both a belief and a believing (which is a doing). One might worry that allowing the possibility of believings damages the distinction that I am trying to make. I don't think so. Even if there are believings (ie, belief-acts), there are also doings that are mere acts, and not belief-acts.

distinction between epistemic rationality and practical rationality. Only beliefs can be epistemically rational; beliefs and doings can both be practically rational.

However, one might wonder about the distinction between practical rationality and morality; after all, the object of both of these domains is an action. Nonetheless, that there is a distinction between practical rationality and morality is easy to see (though it is hard to specify). It is easy to see because one can clearly be practically rational while nonetheless doing something that lacks moral justification, and is perhaps even morally forbidden. The student is not typically morally justified in cheating on the test, even if doing so is a great means to her ends. Similarly, one can be morally permitted to do some act that is ruled out by practical rationality. I am morally permitted to eat ice cream for supper tonight though, given my goal of losing weight, that act is practically irrational for me. Moreover, it makes some sense to say that a mouse is practically justified in taking a certain path through a maze, but no sense to say that that mouse was morally justified in taking the path.

An understanding of the relationship between practical rationality and morality is a key component in seeing why justified closure of inquiry cannot be reduced to purely non-epistemic norms. For this reason, this chapter is largely organized around the different ways these two domains might relate. For example, it might be the case that morality reduces to practical rationality (or vice versa). Moreover, there may be instances in which one lacks a species of rationality in way that does not count as an instance of irrationality. For example, an act could lack practical rationality in a way that does not make the act practically irrational. This latter possibility stems from the seeming possibility that an act could be neither

practically rational not practically irrational and that an act might be neither morally justified nor morally unjustified.

The first complication (that one of these normative domains might reduce to the other) highlights the possibility of a theory of justified closure of inquiry that relies only on practical rationality, and another that relies only on moral permissibility. The second complication (that an act may lack a type of rationality in a way that does not count as an instance of irrationality) makes it possible to articulate theories that combine the two domains of normativity not only through simple conjunction or disjunction, but through some combination of the presence of one kind of justification along with the mere absence of the second kind of justification.

There are at least twenty ways these two domains could relate. Two possibilities are fully reductive; moral justification could reduce to practical rationality (or vice versa). To see how the remaining 18 principles might be articulated, consider that an act might be practically rational, practically irrational, or neither practically rational nor irrational. Also, an act might be morally justified, or morally unjustified, or neither morally unjustified nor morally justified. There are nine ways these possibilities could pair, and each of these nine could become a disjunction or conjunction. (For example, an act might be practically rational *and* morally justified. Alternatively, an act might be practically rational *or* morally justified.) This creates eighteen possible ways these two domains of normativity could relate, in addition to the two fully reductive accounts.

Fortunately, many of these possible combinations can be dismissed in groups. For example, if I can show that an act is always either morally justified or morally unjustified, then I eliminate six of the twenty combinations. If I can show that an act is always either practically rational or irrational, then I am down to ten possibilities. Since no act that is morally unjustified can ever be all-things-considered justified, then I am down to just six possibilities.⁷ It would be exceedingly odd to hold that an act is justified if and only if it is practically irrational and [/or] morally justified. Quickly, I am down to one of the fully reductive possibilities, or to the conjunction or disjunction of the two fully reductive possibilities.

First, I will argue that every act is either practically rational or practically irrational. (Notice that this is really more than I need to show; I really only need to show that every act of closing inquiry is either practically rational or practically irrational). Second, I will show that every act is either morally justified or morally unjustified. Third, I will consider each of the remaining four possibilities in turn. Finally, I will show that we should not be surprised that none of them worked. It turns out that a theory of justified closure of inquiry cannot be reduced to purely non-epistemic normativity.

⁷ Some philosophers follow Bernard Williams in arguing that an act can be morally unjustified and yet all-things-considered justified. Eventually, I will respond to these arguments. For now, I help myself to the *prima facie* plausible view that no act that is morally unjustified can be all-things-considered justified.

Every Act is Either Practically Rational or Practically Irrational and Morally Justified or Immoral

If every act is either practically rational or practically irrational, then we can dismiss a major group of potential theories of justified closure of inquiry that rely on purely non-epistemic normativity. In order to know whether or not it is possible for an act to be neither practically rational nor practically irrational, we need to know what it is for an act to be practically rational. Therefore, I start by developing a case for what it is for an act to be practically rational.

Practical rationality is a species of instrumental rationality in which the object of evaluation is an act. The musician whose end is cutting his ear off is practically rational in choosing the knife over the spoon. It is possible, of course, for the same act to be rational relative to one end and irrational relative to another. Given my end of enjoying my evening, a bowl of ice cream is rational for me. Given my end of losing weight, it is not. It also seems possible for two acts to be equally rational relative to some given end. Eating the chef salad or the house salad may be equally rational; either act is practically rational (though the act of eating both salads may not be).

Given these considerations, here is something we may be tempted to say: once the relevant end is adequately specified for some given agent, an act is either a way of reaching that end, or it is not. If it is, then the act is practically rational. If it is not, then the act is practically irrational. While we may be tempted to say this, it is false.

Some act could achieve an end in a way that one has no reason to expect. It is not rational for me to invest all of my money on a single spin of the roulette wheel

(given that my end is the acquisition of wealth), and this is true even if I win. If I were to win, then the bet would have been an effective means to my end. Such cases make it clear that the fact that an act actually is or is not a good means to an end is not what decides whether an act is practically rational. Here is a plausible explanation for what is wrong with the claim that an act is practically rational relative to some end if it is a means of achieving that end: It fails to account for acts that achieve a given end for an agent, but the agent was unjustified in believing that the act would achieve the end. If this is the correct explanation of the failure of the principle under consideration, that it seems as though the fact that an agent has good reason to believe that an act will achieve one's end is a necessary condition on the act's being practically rational.⁸

That an agent is justified in believing that an act will achieve an end is merely a necessary, and not a sufficient, condition for the act to be practically rational relative to that end. To see this, consider the following case: John has the end of asking the speaker a question. He might raise his hand. He might write out "I'd like to ask you a question" with a black marker on a sheet of paper and hold it up. He might use a megaphone to announce his desire to ask a question. Any of the three alternatives, it is easy to suppose, would suffice to achieve the end he has in mind. But the first option would come with fewer unfortunate side-effects. It would not only achieve the specific end in question, it wouldn't cause others in attendance to think him strange. For an act to be practically rational, it is not enough that the

⁸ I need take no stance on what constitutes a good reason to believe that an act is a suitably good means to an end. This could be fleshed out in terms of expected utility, epistemic probability, reliability, or any number of other ways.

agent has good reason to think that it will achieve the end in question. The agent must have good reason to think that the act must achieve the end in a manner that is at least as good as any other available means to that same end, where the ordering of good means is sensitive to not just the effectiveness of achieving the specific end in question, but also to other ends that one has.⁹

Instead of saying that one needs to be justified in believing that a given course of action is a good means to one's end, we should say that one needs to be justified in believing that a given course of action is as good as any other available alternative. Every belief is either justified, or it is not. The belief that a given act is as good as any other alternative in the pursuit of some end is no different. For every act, the agent is either justified in believing of it that it is as good as any other alternative in achieving the relevant end, or the agent is not so justified.

Suppose an agent is not justified in believing of a given act that it is as good as any other alternative in achieving a relevant end. It follows from this (given what I have said about practical rationality) that the act is not practically rational. It does not yet follow that the act is practically irrational. In order to achieve that result, I also need the claim that an act is practically irrational for one whenever one is unjustified in believing that it is as good as any other alternative in achieving the relevant end.

⁹ Alex Pruss noticed that this articulation is problematic in cases in which there are infinitely many means. While the matter is unclear, I think the means that one has is a psychological fact, and that we can't have infinitely many means. If the reader is concerned about this problem, Pruss helpfully suggested that "as good" could be changed to "close enough to as good."

Such a claim is plausible, but faces a challenge. Suppose one finds oneself in a situation such that one has a certain end, but is not justified in believing that any single means to that end is as good as any other alternative. Are all pursuits towards that end practically irrational? To the extent that this seems a bad result, it counts against the conclusion I need. But this result is not really so bad. There are many instances where our evidence is insufficient to think that any single means is the best means; but one need not have evidence sufficient to justify the belief that the means is the *best* means. One need only have reason to think that the means is *as good as any other* means. The extent to which one lacks evidence to justify one in thinking some given option is the best option is the extent to which one's field of *as good as any other* options increases. For example, a young man wooing a girl he just met might be considering where he should take her to eat. To the extent that he lacks reason to believe of any given option that it is the best option, he has reason to believe of some given option that it is as good as any other option. He may have no idea what kind of food she likes. That only makes it more likely, for him, that some given restaurant is as good as any other. In other words, it is not obvious that one ever lacks justification for believing of some means that that means is as good as any other means.

I should also note that I do not need to argue that being justified in believing of some act that it is as good as any other act in the achievement of some end is sufficient for the act to be practically rational; I need only hold that it is necessary. I also need to hold that an act is practically irrational relative to some end whenever one is not justified in believing that it is as good as any other alternative.

There is an objection to the thesis that an act is not practically rational for one unless one is justified in believing that the act is as good as any other act in the pursuit of the relevant means. The objection is this: one need not be *justified* in believing that the act is as good as any other; it is enough that one unjustifiably believe that the act is as good as any other. Suppose I have the end of acquiring wealth, and I believe that investing all of my money on a spin of the roulette wheel is as good as any other means in the pursuit of that end. As such, one may hold, the act of betting it all would be practically rational for me.

The problem with such a view is that it is at odds with the fact that practical rationality is an achievement. If practical rationality could be “achieved” merely by coming to a certain unjustified belief, then hypnosis could make nearly any act practically rational. If such a claim were true, practical rationality would not be understood as a success; but it is understood as a success.

If my thesis about practical rationality is correct, then there can be no account of justified closure of inquiry that reduces the normative domain that governs legitimate closures solely to the realm of practical rationality, since practical rationality itself is dependent on epistemic rationality. This creates confusion for the next section of this chapter, in which I consider a principle that attempts to do just that. Nonetheless, I intend to consider the plausibility of reducing closure of inquiry to practical rationality. If my thesis about practical rationality is correct, it won’t truly be a full reduction. However, it might be that my thesis is mistaken; besides, even if my thesis is correct, I still need to consider the

principle in question, even if it is not accurately described as a fully reductive principle.

Before dealing with the fully reductive theories, however, I still need to show that every act is either morally justified or morally unjustified. It seems likely that every act is either a breach of one's moral obligations or not. If the act is a breach of one's moral obligations, then the act is morally unjustified. If the act is not a breach of one's moral obligations, then the act is morally justified.

This is not to say that it is never possible for acts to be better or worse than others, morally speaking. I need take no stance on that. All that I need to secure my result is to show that every act is either a breach of moral obligations or not.

If one thinks that moral dilemmas are possible, and that in such cases one is morally justified in breaching a moral obligation, then (perhaps) we get a case in which an act is both morally justified and morally unjustified. While I think it unlikely that there are any moral dilemmas, I need not argue as much. Even if moral dilemmas were possible, and the above language is the correct way to describe the moral situation, it still is not a problem for the claim that every act is morally justified or morally unjustified.

Since we have good reason to think that every act is either practically rational or practically irrational and morally justified or immoral, then we have good reason to reject a large group of possible ways in which practical rationality and morality might combine to form a principle of justified closure of inquiry. I now focus on each of the four remaining possibilities.

The First Fully Reductive Possibility: Reduction to Practical Rationality

Two of the four remaining ways that practical rationality and moral justification might combine to express an adequate theory of justified closure of inquiry are fully reductive; that is, they reduce an explanation of when closure is justified down to a single domain of normativity. First, I will deal with a fully reductive theory that reduces the normative domain to the practically rational.

PR: S is justified in closing inquiry into whether or not p iff it is practically rational for S to close inquiry into whether or not p .

The prima facie plausibility of principles like PR stem from the close connection of practical rationality to intentional acts. Since I've defined a closure of inquiry as an act, it is prima facie plausible that practical rationality is what governs it.

The obvious objection to the view is that morality often constrains practical rationality. An act that is practically rational and morally unjustified for one is not an act that one is all-things-considered justified in doing. (In what follows, I'll call this the obvious objection). If PR can be defended from the obvious objection, then the simplicity of PR makes PR highly appealing. There are at least two ways a defender of PR might respond. In the next two sections, I will consider these two responses. I argue that neither response works, and, as such, no theory of justified closure of inquiry that makes use of only practical rationality will be successful.

Denying Moral Normativity

The most interesting and important of the two possible replies is this: One might claim that morality can never clash with and trump practical rationality. There are at least two ways one might go about defending this claim. First, one

might argue that all propositions about what one morally ought to do are false.

Second, one might argue that no claims about what one morally ought to do express propositions, and hence none of them are true.

The first claim is connected to error theory, and the second is connected to expressivism. I should hasten to add that expressivists need not think that morality can never clash with practical rationality. But it is possible for one to make use of the claim that 'S ought to φ ' does not express a proposition as a means of making the case that morality can never clash with practical rationality. After all, if there are no true propositions about what one morally ought to do, then no true propositions about what one morally ought to do clash with true propositions about what one practically ought to do.

Whether my opponent holds that 'S ought to φ ' does not express a proposition, or that 'S ought to φ ' is always false, the crucial claim in this response is that my opponent must hold that morality can never clash with practical rationality. Since this claim is widely contradicted by common pre-theoretical belief, my opponent needs an argument.

Arguments in favor of error theory, of which John Mackie has been the leading proponent, might easily be appropriated to show that 'S ought to φ ' is always false.¹⁰ The Stanford Encyclopedia of Philosophy entry on moral anti-realism says that error theorists think that "when we say 'stealing is wrong' we are

¹⁰ Mackie, *Ethics*; Mackie, *The Miracle of Theism*.

asserting that the act of stealing instantiates the property of wrongness, but in fact nothing instantiates this property, and thus the utterance is untrue.”¹¹

It is extremely important to note, however, that it is not enough merely to argue that there are no moral *properties* in order to make this response to my objection to PR. Judith Jarvis Thomson believes there are no moral properties, but one who accepts her position could not save PR.¹² In order to save PR from the objection, one would need to claim not only that there are no moral properties, but also that there are no moral facts. They need to claim that all statements like “S ought to φ ” are either false or express something other than a proposition.

Thomson, though she agrees with non-cognitivists and error theorists that there are no moral properties, thinks that statements like ‘S ought to φ ’ are made true on the basis of what she calls “goodness fixing kinds.”

For Thomson, S ought to φ just in case S belongs to a kind K that admits of defect, and S would be a defective K if S did not φ .¹³ There are, of course, a multitude of qualifications in Thomson’s theory but the basic idea is that morality dictates that the judge not release the innocent man, even if doing so would help quell the riots, if doing so would cause the judge to be (or show that the judge was) defective qua person (or, perhaps, qua judge). Thus, while Thomson agrees that there is no moral

¹¹ Joyce, “Moral Anti-Realism.” Error theorists differ in their handling of statements like “S ought to φ .” Those who think that such statements can express true propositions, and even in those cases in which the ought in question is an ethical ought, can not save IR.

¹² Thomson, *Normativity*.

¹³ Jay Wallace, “‘Ought’, Reasons, and Vice.”

property (such as goodness), she does think there is the property of being a defective kind, and from the concept of a defect relative to a kind she develops an account of moral normativity. If Thomson is right, then there are no moral properties, but the judge still ought not do what instrumental rationality recommends that he do. If such conflict between morality and instrumental rationality is possible, then PR is doomed.

My opponent needs to do more than merely argue that there are no moral properties (perhaps on the basis of their purported queerness); she needs to argue that statements of the form 'S ought to φ ' are either always false (at least when the ought in question is a moral ought) or fail to express a proposition. Error theorists have argued that there are no moral properties, but, even granting that there are no moral properties, there is no good reason to think all ethical statements are false.

While error theorists think ethical statements always express false propositions, expressivists claim that they express no proposition at all. If this is true, no statement of the form 'S ought to φ ' expresses a true proposition. If that were the case, it would be reasonable to suppose that there were no moral facts. Expressivists think that such statements express mental states.¹⁴ Typically, expressivists are also non-cognitivists. A non-cognitivist is one who thinks that statements like 'S ought to φ ' express a non-cognitive mental state (like a desire or emotion) rather than a cognitive mental state (like a belief).

A.J. Ayer, an early expressivist, held that ethical language did not contain "factual content," but simply showed that the speaker attended the concept with a

¹⁴ Blackburn, *Essays in Quasi-realism*; Gibbard, *Wise Choise, Apt Feelings*.

certain feeling of approbation or approval. According to Ayer, the function of ethical words is to “express feelings about certain objects, but not to make any assertion about them.”¹⁵ An early objection to Ayer’s emotivism was that it entailed that moral disagreement is impossible.

When I say “stealing is bad” I am not, according to the expressivist, uttering a proposition. I am merely making known a fact about my mental states. According to emotivists, I am saying something like “stealing...booo!” When my interlocuter then replies, “stealing is not bad,” they are also simply expressing their feeling about stealing. As such, there is no disagreement. Emotivists did not think this worry was a problem for their view; they simply insisted that this is in fact the case.

The problem of moral disagreement, which many emotivists could accept, eventually led to an even more serious problem. The problem of embedding, first discussed by Geach (and inspired by Frege), is a further stage in the worry about the fact that people seem to be able to disagree about ethical matters. Geach pointed out that ethical judgments can be embedded in other utterances about which it is implausible to claim there is not propositional content. I might say “If stealing is wrong, then downloading illegal music off the internet is also wrong,” or “Mark believes stealing is wrong,” or “Either stealing is wrong or it is not wrong.” In none of these three sentences do I appear to be straightforwardly expressing an emotion about stealing, and yet the truth-value of each of them is dependent upon the embedded ethical claim. Many people came to think that the problem of embedding was a decisive claim against emotivism. However, contemporary expressivists

¹⁵ Ayer, *Language, Truth and Logic*. p 107

believe that more complex theories can handle the problem of embedding without being forced to accept factualism about ethical judgments. Alan Gibbard is perhaps the most influential of these.

While Ayer eliminated ethical normativity by reducing ethical facts to emotions, Gibbard eliminates ethical normativity by reducing ethical facts to a special kind of mental state, a mental state he calls a normative judgment. A normative judgment, Gibbard claims, is not an emotion, but a judgment that has the fundamental quality of “disagreeing with” other mental states and actions. “According to Gibbard,” writes Wedgwood, “the content of the normative judgment is determined by the actions and other mental states it disagrees with—and especially by the *plans* and *attitudes of disagreements* that it disagrees with” (45). Plans might be contingent (what one would do in some circumstance in which one might, or might not, find oneself) or hypothetical (what one would do in any situation, even if there is no chance one would find oneself in that situation).

From these concepts, Gibbard explains how some plans might be complete and consistent. He calls such plans *hyperplans*. A plan is complete if it contains a plan to do one of the alternative actions that are available relative to any possible circumstance and to have or not have every mental state that is available in that circumstance. This includes a plan to either *forbid* (disagree with) a given mental state or action or *permit* (disagree with disagreeing with) a given mental state or action. A plan is consistent if it contains no logically inconsistent beliefs, is realizable (that is, it is logically possible to realize that circumstance), and, for every

mental state the plan contains, it will never also contain a plan to disagree with that mental state. If a plan meets these criteria, it is a hyperplan.

With this framework in place, Gibbard can solve the problem of embedding. The meaning of the phrase 'stealing is wrong' is identified with "an attitude of disagreeing with an action or mental state, or with an attitude of permitting (that is, disagreeing with disagreeing with) an action or mental state."¹⁶ We can then explain the content of ethical statements in terms of the hyperstates they allow. If a hyperstate contains a plan to disagree with an ethical judgment, then that ethical judgment does not allow that hyperstate. Embedded ethical statements, whether they are embedded in disjuncts, conditionals, or propositions about someone's beliefs can be explained similarly, following ordinary logical procedures. For example, we can explain the content of the sentence *either stealing is wrong or it is not wrong* as itself a judgment constituted by the disjunct of two atomic judgments. This disjunctive allows a hyperstate if and only if at least one of its disjuncts allows the hyperstate. Since a hyperstate must be complete and consistent, every hyperstate will allow this disjunctive in virtue of the fact that every hyperstate must allow one of the disjuncts. Gibbard used this framework to explain the content of propositions that contain embedded ethical judgments.

The problems for Gibbard are two-fold. First it is not clear that mental states are the kinds of things that can disagree, at least not if our notion of disagreement is psychological.¹⁷ Attitudes can be logically inconsistent with one another, but they

¹⁶ Wedgwood, *The Nature of Normativity*, 46.

¹⁷ Kvanvig, "Certain Doubts."

can't disagree, at least not in any psychological sense of disagree. People disagree with one another in virtue of the attitudes that they have, but the attitudes themselves don't disagree. If we are going to understand what Gibbard means by his talk of disagreement, then, we are going to need to search for some other notion. Since the presence or absence of disagreement seems to generate, or leave absent, justification for having certain mental states or plans or plans to have certain mental states, we need some notion of disagreement that could itself explain that relation between the having of justification and the presence or absence of disagreement. Further, the task of the expressivist is to explain ethical normativity in strictly psychological terms. If disagreement, perhaps the most fundamental element in Gibbard's theory, is inexplicable in such terms, it is no longer clear that the utterances in question are being reduced to mental states.

Wedgwood articulated a second, slightly more complicated, problem for Gibbard. He noticed that any satisfactory explanation of why it is bad to make an inconsistent set of judgments commits one to factualism. In order to explain Wedgwood's objection more fully, it is necessary to say something about a connection between the normative evaluation of statements and the normative evaluations of judgments. Statements are, of course, evaluated in a variety of ways, as rude, funny, or inappropriate, for example. But statements are also evaluated in a special way simply *as statements*. Wedgwood is not as clear as one might like about what it means for a statement to be evaluated simply as a statement. I think he has in mind the following kind of thing: suppose during a lonely walk one night I look up at the moon and say, "The moon is made of green cheese" or "murder is not really so

bad.” There is a sense in which these statements are unjustified, even though their failure as statements can’t be linked to social expectations or similar evaluative criteria. Their failures as statements are closely linked to the failures of the judgments to which they correspond. If a statement is justified according to this special kind of evaluation, then the judgment to which it corresponds should also be justified. This close connection between statements and judgments gives us good reason to think that “an explanation of the meaning of a normative statement must provide some account of the conditions or standards of justification or warrantedness” (50).

Gibbard thinks that his account of the meaning of ethical statements can provide such an account. He thinks that inconsistency in regards to one’s normative judgments, both in regards to one’s own judgments and the judgments of one’s community, is bad. The problem is in explaining why, according to his theory, we should think that such inconsistency is bad. Wedgwood helpfully quotes Gibbard as follows:

A set of judgments is inconsistent if there is a hyperstate that every judgment in that set allows. It is consistent otherwise: it is consistent if every possible hyperstate is ruled out by one or another of the judgments in the set. If, then, my judgments are inconsistent, there is no way I could become fully opinionated factually and fully decided on a plan for living—no way that I haven’t, with my judgments, already ruled out (59).

The problem with this approach is that it is hard to imagine what is so bad about the fact that my judgment is ruled out by all possible hyperstates. Hyperstates are ideals anyway and, as Gibbard concedes, none of us will ever be in one. As such, explaining what is bad about having inconsistent judgments by pointing out that

doing so rules out all possible hyperstates is unsatisfactory; ruling out all possible hyperstates isn't so bad. The response is supposed to explain what seems bad about asserting a seemingly false ethical statement. It does so by claiming that the statements in question exhibit inconsistent judgments, but then explains inconsistent judgments in a way that doesn't seem bad. As such, it is hard to understand how it is an explanation for that which it attempts to explain.

Wedgwood goes on to give a convincing case that any possible account of what is bad about inconsistent judgments commits one to factualism. Recall that the evaluation of a statement *as a statement* should mirror the evaluation of the corresponding judgment. Wedgwood plausibly supposes that there is some desirable property in positively evaluated statements that is also present in (or corresponds to some desirable property present in) positively evaluated judgments. If a statement is evaluated positively as a statement then the judgment with which the statement corresponds will also be evaluated positively as a judgment. He then argues that we know enough about the property that is found in positively evaluated statements/judgments to suppose that it simply is *truth*. He suggests four things we know about this property which, taken together, strongly suggest that the property is truth.

First, the property is such that no two logically inconsistent statements can both instantiate that property. Second, any statement that can be inferred from other statements that have this property must also have this property. Third, statements with this property follow the ordinary rules of logic when it comes to their relations insofar as this property is concerned; for example, a statement with a

conjunction has this property only if each of the conjuncts has the property. Finally, if one accepts that some given sentence has this property, then one is also committed to accepting the statement that the sentence expresses.

To summarize, any attempt to save PR by holding that all statements of the form ‘S ought to φ ’ are either false or fail to express propositions is unsatisfactory. Regarding the first disjunct, error theorists have given arguments that there are no moral properties. However, Thomson has shown us that this is not enough. They also need to give us an argument that all moral propositions, especially of the kind ‘S ought to φ ’, are false. Regarding the second disjunct, attempts by expressivists to sidestep the problem of embedding leave them working with concepts that are not psychological. Further, any plausible attempt to explain what is bad about inconsistent judgments forces them to suppose a property that is found in “good” judgments that is very likely the property of being true.

*Morality as a Species of Instrumental Rationality, or Morality as a Constraint
on the Permissible Ends of Instrumental Rationality*

The arguments of the proceeding section show that neither error theory nor expressivism give us the tools with which to plausibly deny that there is any such thing as moral normativity. However, a less radical reply might save PR. Namely, one might claim that moral normativity is a species of practical rationality.

It seems clear, however, that morality is not a species of practical rationality. If morality is a species of practical rationality, then one can never, while acting in a morally permissible manner, simultaneously, in regards to the same act, fail to be practically rational. But one can be practically irrational while acting in a morally

permissible manner. Suppose that I decide to send \$100 to UNICEF. I prefer to do so as quickly as possible. I could send the \$100 by mailing a check or by using a bankcard. Suppose that, despite my sincere preferences for the results that would ensue upon my using the bankcard (perhaps because I don't have the time to write out a check), I send a check. Sending the check is clearly practically irrational, and yet surely it is morally permissible. The seeming possibility of cases like this gives one who wants to hold that the moral is a species of the practically rational an uphill battle.

The Second Fully Reductive Possibility: Reduction to Morality

It might be helpful to take a moment and recall the flow of our investigation. Since morality and practical rationality are often thought to be the domains of normativity that govern the justificatory status of actions, and since closure of inquiry is an action, I am investigating the logical space available for a theory of justified closure of inquiry that relies only on these two domains of normativity. I identified twenty possibilities, but quickly limited the field to just four plausible candidates. Two of those four candidates are fully reductive; that is, they appeal to only a single domain of normativity to explain the justificatory status of a closure. I considered and rejected one of those two fully reductive options in the last major section. In this section, I will consider the second fully reductive possibility. Ultimately, I reject it as well.

A theory of justified closure on inquiry that relies solely on moral normativity to explain what justifies a closure might claim the following:

MJ S is justified in closing inquiry concerning whether or not p iff S is morally justified in closing inquiry into whether or not p .

If someone believes that practical rationality reduces to morality (perhaps because all instrumental rationality so reduces), it would be natural for one to accept MJ or some similar principle. A defense of MJ, then, would be bolstered by a defense of the claim that practical rationality reduces to morality. By modifying an argument formulated by Trent Dougherty in defense of the claim that each instance of epistemic irresponsibility/irrationality is just an instance of non-epistemic irresponsibility/irrationality, we can formulate a strong *prima facie* argument in favor of the reduction necessary for MJ.¹⁸ I'll call this argument *The Reduction*:

5. If practical rationality can be reduced to a more fundamental normative concept without losing the ability to explain relevant data, then the reduction ought to be accepted.
6. Morality is more fundamental than practical rationality.
7. Practical rationality can be reduced to morality without losing the ability to explain the relevant data.
8. The reduction ought to be accepted

(1) seems to follow from plausible assumptions regarding the value of simplicity in theory, and I find (2) highly intuitive. In any case, I take the truth of (2) for granted for our present purposes. What about (3)? It is neither obviously true nor obviously false. How might we go about investigating it?

I need first to say what is meant by 'relevant data'. In deciding what sorts of intentional acts are or are not justified, we begin by looking at obvious cases in which someone is or is not justified in acting in a certain way. By 'relevant data' I mean those obvious, individual cases of, for example, blameworthiness. Can we explain (to take an example of an individual judgment) why the judge is

¹⁸ Dougherty, "Reducing Responsibility."

blameworthy when she sentences the innocent man to die as a means of achieving the end of quelling the riots, and can we do so without appealing to a failure of practical rationality? After examining a broad range of such cases, can we in every case give such an explanation? If so, then perhaps we have good inductive support for (3).

In some cases, it certainly seems like we can explain the relevant data without appealing to a failure of practical rationality. The problem is that, for The Reduction to be successful, a similar explanation must be available for all relevant data. A single counterexample is sufficient to show the falsity of (3). There are such counterexamples.

Suppose that Lucy, who loves lollipops, has the desire to eat a lollipop. Lucy is a virtuous, moderate person, who wisely spends money. She has plenty of money in her pocket, plenty of time on her hands, and she knows that she can buy a lollipop in the atrium of the skyscraper she happens to be standing in front of. Unfortunately, Lucy has a severe psychological aversion to entering skyscrapers on Thursdays, and today is a Thursday. She knows full well that, if she were to enter the skyscraper, her feelings of anxiety would immediately fade. She would not only be proud that she entered the skyscraper, but she would be in a position to buy the lollipop that she desires. Nonetheless, she is unwilling to enter the skyscraper. There is a sense in which Lucy has failed in terms of what rationality dictates that she ought to do, but that failure cannot be explained by any moral failure on her part. It seems, then, that practical rationality cannot be reduced to morality, and the argument proffered by the proponent of MJ is unsuccessful.

There are other compelling reasons to think The Reduction must be unsound, because there are good reasons to think MJ must be false. MJ entails that anyone is justified in closing inquiry anytime they are morally permitted in closing inquiry. But this seems clearly false. One can be morally permitted to close inquiry even when it is not practically rational for one to close inquiry, and in such cases one is not justified in closing inquiry. Suppose that my goals and desires make it practically rational for me to eat at a new restaurant in town. Suppose also that I don't know whether the restaurant is downtown or uptown. I am not morally forbidden from closing inquiry into <the restaurant is downtown>, but neither am I justified in closing inquiry.

In the last section, I showed that the first of two possible fully reductive theories of justified closure of inquiry is implausible. In this section, I showed that the second of two possible fully reductive theories of justified closure of inquiry is implausible. I will now move on to the final two possible combinations of moral and practical justification that could provide a fully non-epistemic explanation of justified closure of inquiry. The first possibility is the conjunction of practical rationality and moral justification. The second is their disjunction. If the principle that explains when a closure of inquiry is justified is expressed as a complex sentence constructed from the presence or absence of atomic sentences that are themselves made true by the presence or absence of moral justification or practical rationality, then those atomics must be constructed from truth-functional connectives. Hence, all we have left to investigate is the conjunction and disjunction of MJ and PR.

The Conjunction

I have considered arguments in favor of PR and MJ, and have concluded that neither of them succeeds. The argument I considered in favor of PR fails because practical rationality is not the only normative concept that regulates justified acts. The argument I considered in favor of MJ fails because practical rationality cannot be reduced to morality.

This chapter has been driven by some fairly common philosophical assumptions that I don't yet want to abandon. I have assumed that there are exactly three relevant normative domains - the practical, the moral, and the epistemic. The practical is governed by practical rationality, the moral is governed by moral norms, and the epistemic is governed by epistemic normativity. I have assumed that intentional acts are not governed by epistemic normativity. I have assumed, for our present purposes, that closings of inquiry are intentional acts, and that beliefs are not intentional acts. If all of these assumptions are correct, it follows that closure of inquiry is governed solely by moral norms, solely by instrumental rationality, or by both. I have argued that it is not governed solely by moral norms and it is not governed solely by practical rationality. Perhaps it is governed by a combination of the two? Consider the following principle:

PR \wedge MJ: S is justified in closing inquiry concerning whether or not *p* iff it is practically rational for S to close inquiry into whether or not *p* and S is morally justified in closing inquiry into whether or not *p*.

A first concern for PR \wedge MJ is this: one might be morally obligated to close inquiry even though such closure is practically irrational. If that is possible, then there are cases in which PR \wedge MJ fails because S is justified in closing inquiry in some

instances in which it is not practically rational for S to close inquiry. Since the fact one is morally obligated to perform an act is sufficient to show that one all-things-considered ought to perform that act, the proponent of $PR \wedge MJ$ needs to hold that such situations are impossible. But such situations seem possible. The easiest fix is to add a disjunct to the principle in question:

$(PR \wedge MJ) \vee MO$: S is justified in closing inquiry concerning whether or not p iff
 a) it is practically rational for S to close inquiry into whether or
 not p and S is morally justified in closing inquiry into whether
 or not p , or b) S is morally obligated to close inquiry.

However, $(PR \wedge MJ) \vee MO$ faces problems of its own; in order to be truly reductive of epistemic normativity, it must be accompanied with a theory of practical rationality that does not itself appeal to epistemic rationality. Recall my earlier articulation of what it is for an act to be practically rational for one: one has to be justified in believing that a given course of action is as good as any other available alternative.¹⁹ Given my account of practical rationality, $(PR \wedge MJ) \vee MO$ is not reductive. One who wants to eliminate epistemic normativity from a principle of justified closure of inquiry obviously needs a different account of practical rationality.

Some philosophers, such as I.J. Good, have argued that practically rational agents maximize expected utility.²⁰ When deciding whether or not to perform some act, the rational agent will update her beliefs given all of her available evidence.

¹⁹ In chapter three, I will make an important distinction between two kinds of truth-oriented justification, and this distinction will further elucidate this articulation of practical rationality.

²⁰ Good, "On The Principle Of Total Evidence."

Then, she will do the act that maximizes her expected utility. Suppose she is trying to decide whether or not she should have a surgical procedure. She should multiply her degree of belief that a certain outcome will result by the utility (for her) of the outcome in question. She should sum the expected utility of doing the act, and she should sum the expected utility of not doing the act. She should do the act if the former number is greater than the latter.

Under this conception of practical rationality, $(PR \wedge MJ) \vee MO$ says that one is justified in closing inquiry just in case doing so maximizes one's expected utility and one is morally justified in closing inquiry (or one is morally obligated to do so). Good showed that it always maximizes expected utility to further inquire, provided that the inquiry is cost free.

Nonetheless, there is strong evidence that the practically rational agent need not always maximize expected utility. Lara Buchak argues that risk-avoidant agents are justified in closing inquiry (even when the inquiry is cost free) if the agent's concern that the experiment will result in misleading evidence is sufficient to increase her subjective probability of a worst state of affairs, even if doing the experiment also increases the probability of the best state of affairs.²¹ She gives the example of a ship captain who, standing on the shore, sees nine people sinking in rough waters. The captain has just enough time to perform a test on her boat that will give her more information regarding the sea-worthiness of the boat. She is fairly confident that her boat will handle the rough waters and is concerned that the

²¹ Buchak, "Instrumental Rationality, Epistemic Rationality, and Evidence-Gathering."

test will provide misleading evidence. If the test provides information that makes it less likely that the boat is sea-worthy, she knows she will lose her nerve and not attempt to save the nine sinking souls, who will surely die if she does not save them. Buchak argues that it is practically rational for risk-avoidant agents to decline to perform the test. Performing the test brings with it the risk of the state of affairs that would result from misleading evidence. For utility maximizers, that risk is always offset by the possibility of getting non-misleading evidence. For risk avoidant agents, the potential gain needs to be greater than it needs to be for non-risk avoidant agents.

Some may insist that risk-avoidant agents are not practically rational. Nonetheless, there is strong intuitive evidence that risk-avoidant agents are practically rational. Allais' paradox offers support that such agents are rational.²² People are asked which of the following two pairs of gambles they would select, if given the choice:

- L1: \$5,000,000 with probability 0.1, \$0 with probability 0.9.
- L2: \$1,000,000 with probability 0.11, \$0 with probability 0.89.

and

- L3: \$1,000,000 with probability 0.89, \$5,000,000 with probability 0.1, \$0 with probability 0.01
- L4: \$1,000,000 with probability 1.

Most people choose L1 over L2 and L4 over L3, despite the fact that there is no way to assign utility to dollar amounts such that L1 has a higher utility than L2 and L4 has a higher utility than L3. Many people continue in this pair of choices even after being shown that there is no assignment of utilities to dollars such that this pair of

²² Allais, "Criticisms of the Postulates and Axioms of the American School."

choices maximizes expected utility. When an act is called irrational, we mean (at least) that a fully informed and virtuous or well-formed agent would not choose to do the act. There is nothing about those people who continue to choose L1 over L2 and L4 over L3 that appears vicious or ill-formed. Further, the fact such people can often fully explain the fact that there is no way to assign utility to dollar amounts such that L1 has a higher utility than L2 and L4 has a higher utility than L3 makes it seem as though their choice is not a failure of information or comprehension.

If an act is practically rational for an agent if and only if it maximizes the expected utility for that agent, then risk avoidant agents are not being practically rational. But it is not at all obvious that risk avoidant agents are failing to be practically rational.

Furthermore, maximization of expected utility accounts of practical rationality demand that probability and utilities are sufficiently fine-grained for such procedures to be helpful; but it seems to me that they are not. Our subjective beliefs and subjective utility assignments are far too coarsely grained to be helpful in such calculations. Further, there is also the question of regress of calculations. If my closing an inquiry is only practically rational if it maximizes expected utility, then I can only know if my closing inquiry is practically rational by performing the calculations. But performing the calculations is itself an inquiry – a difficult and complex inquiry. In order to know whether or not I should close the inquiry that is itself required in order for me to know if I am practically rational in closing inquiry, I must perform yet another inquiry. A regress is launched.

Some will not be convinced by these concerns that practical rationality can't be reduced to the maximization of expected utility. For those who remain convinced that practical rationality is about utility calculations, I recommend the second of my two arguments in this chapter; it is very brief, and I will present it in the last section of this chapter.

The Disjunction

The final of the four plausible ways in which practical rationality and moral justification might combine to yield a theory of justified closure of inquiry is the disjunction of PR and MJ.

PR \vee MJ S is justified in closing inquiry concerning whether or not p iff at least one of the following is true: *a*) it is practically rational for S to close inquiry into whether or not p or *b*) S is morally justified in closing inquiry.

PR \vee MJ gets some cases right that PR \wedge MJ gets wrong. One of the problems with PR \wedge MJ is that it fails in those cases when one is morally required to close inquiry even though it is not instrumentally rational for one to do so. PR \vee MJ does not fail in this way; one is obviously morally justified in closing inquiry on those occasions when one is morally required to do.

Nonetheless, PR \vee MJ is false. PR \vee MJ fails for the same reasons that the left-to-right entailment of MJ fails; there are cases in which one is not morally forbidden from closing inquiry, yet one is not all-things-considered justified in closing inquiry. Such cases are easily produced. Suppose that my goals and desires direct me to eating at a new restaurant in town. Suppose also that I don't know whether the restaurant is downtown or uptown. I am not morally forbidden from closing inquiry

into <the restaurant is downtown>, but neither am I justified in closing inquiry. But PRvMJ says that I am justified in closing inquiry in such a case, because the second disjunct is satisfied.

Epistemic Normativity and Justified Closure of Inquiry

We should not be surprised that no combination of the presence or absence of practical justification and/or moral justification is able to provide an adequate theory of justified closure of inquiry. After all, there is good reason to think we can hold those two domains steady for an agent, adjusting only the agent's epistemic status relative to the claim into which one closes inquiry, and get differing correct answers to the question of whether or not the agent is justified in closing inquiry. That should not be possible if justified closure of inquiry really was all about the non-epistemic normative domains.

One such case in which one can hold the practical and moral domains steady, adjusting only the epistemic, and yet get differing returns on the question of whether or not the agent is justified in closing inquiry is what I will call the frozen pond case, first discussed by Jeremy Fantl and Matthew McGrath.²³ Imagine you are deciding whether or not to cross a frozen pond. You need to get to the other side, and it will take a long time to walk around the pond. In the first case, you do not know that the ice will not break. In the second, you do know that the ice will not break. In the second case, but not the first, you are justified in closing inquiry into whether or not the ice will break if you walk across it. Imagine if you were to say, "I

²³ Fantl and McGrath, *Knowledge in an Uncertain World*, 73.

know the ice will not break, but I'm going to continue to investigate whether or not the ice will break." People would look at you funny. If you said, "I know the ice won't break, but I'm going to walk miles out of my way anyway," your travel companion would likely and sensibly reply, "Then you don't know that the ice won't break!" If justified closure of inquiry was all about practical rationality and morality, it should not be possible to hold the non-epistemic domains steady, adjust only the epistemic domain, and get a different correct verdict on whether or not the closure was justified. But, as the pond case shows, it is possible. In the next chapter, I will argue in more detail for the conclusion that knowledge justifies closure of inquiry.

I conclude that justified closure of inquiry is not solely regulated by non-epistemic normativity. Henceforth, I will only consider theories of closure that make use of epistemic rationality. In the next chapter, I will consider a purely epistemic principle of closure of inquiry that relies only on synchronic epistemic justification.

CHAPTER THREE

Closure of Inquiry and The Sufficiency of Knowledge

In chapter two, I argued that justified closure of inquiry is not regulated solely by non-epistemic normativity. Such a conclusion is a step in the right direction toward a workable theory of justified closure of inquiry. Given the conclusion of chapter two, we now know that any plausible theory must include a role for epistemic normativity. In this chapter, I am going to present an argument for a principle that further highlights the role of epistemic normativity in justified closure of inquiry. I am going to argue that if S knows p , then S is evidentially justified in closing inquiry concerning whether or not p . I call this claim *The Sufficiency Premise* (abbreviated: TSP), and it is the first step in a larger argument regarding the relationship between epistemic justification and closure of inquiry.

Ultimately, I intend to argue for a nonreductive theory of justified closure of inquiry. That S knows p is *not* a sufficient condition for it to be the case that S is all-things-considered justified in closing inquiry concerning p because, for example, it might be morally reprehensible for S to close inquiry into p . For now, however, I intend to bracket the concern that one might be unjustified in closing inquiry because of the immorality of doing so (or because of some other normative constraint that is unrelated to the strength of one's epistemic position), despite the fact that one has come to know that into which one is inquiring.

My reasons for wanting to temporarily bracket off the concerns of practical rationality and morality are made plain if one considers a principle of justified closure of inquiry that speaks of justified closure only in terms of what epistemic rationality recommends. Consider the following principle:

EJ: S is justified in closing inquiry concerning whether or not p iff S is epistemically justified in believing p or in believing $\sim p$.

Another way to talk about the insensitivity of principles like EJ to the concerns of practical rationality and morality is to say that EJ is insensitive to the distinction between evidentially based closings and closings that are justified on practical or moral grounds. Distinguishing between evidentially based closings and non-evidentially based closings is helpful (and I will make use of such language in this chapter), but deceptive if not understood as shorthand. It is not at all obvious that evidence can ever justify an action. It is plausible to suppose that evidence can and does justify, but the kind of justification it provides relates to doxastic states, not acts. Chisholm, for example, held that only doxastic states, and not actions, were subject to epistemic appraisal. A doxastic state, under Chisholm's view, is either a state of believing, or the "withholding of a believing", or the "preferring of one believing to another believing, or of a believing to a withholding, or of a withholding to a believing, or of a withholding to a withholding."¹ If Chisholm is correct that only states of believings and the like (and not actions) are subject to epistemic appraisal, then it is not at all clear that evidence can ever justify in an epistemic sense.

¹ Chisholm, "Firth and the Ethics of Belief," 123.

Nonetheless, I will talk of S's evidence justifying S in closing inquiry. What I mean by my use of such language is that, *when concerns of non-epistemic normativity are bracketed off*, one is all things considered justified in closing inquiry in all cases in which one is epistemically justified in believing p or in believing $\sim p$. We might modify as follows to make use of this convenient shorthand:

EJ' S's evidence justifies S in closing inquiry concerning whether or not p
iff S is epistemically justified in believing p or in believing $\sim p$.

Notice, however, that EJ' faces another problem. Suppose S knows p . Then, on some influential theories of knowledge², S has undefeated, justified, true belief that p . Notice, however, that one can have epistemic justification to believe propositions that one does not believe. Suppose, then, that S is epistemically justified in believing p , but does not believe p . It follows from EJ' that S is evidentially justified in closing inquiry concerning p , even though S may believe neither p nor $\sim p$. In light of this worry, we might modify EJ' as follows:

EJ'' S's evidence justifies S in closing inquiry concerning whether or not p
iff S's belief that p (or S's belief that $\sim p$) is epistemically justified.

Take this brief introduction to EJ'' as a first step towards its defense. EJ'' is not *prima facie* implausible. In chapters two and three, I defend EJ''. I argue as follows (I'll call this *The Main Argument*):

1. If S knows p , then S is evidentially justified in closing inquiry concerning whether or not p . [premise]
2. Either S's knowledge that p or a fundamental constituent of S's knowledge that p explains why S is evidentially justified in closing inquiry concerning whether or not p . [premise]

² Lehrer and Paxson, "Knowledge," 225. The issue I am discussing here is not unique to defeasibility theories. Under most theories of knowledge, one can be epistemically justified in believing some p even when one does not actually believe p .

3. S's knowledge that p can't explain why S is evidentially justified in closing inquiry concerning whether or not p . [premise]
4. A fundamental constituent of S's knowledge that p explains why S is evidentially justified in closing inquiry concerning whether or not p . [from (2) and (3)]
5. The fundamental constituents of S's knowledge that p are the following: p is true, S has epistemically justified belief that p , and S's belief that p is undefeated. [premise]
6. The feature of S's knowledge of p that explains why S is evidentially justified in closing inquiry concerning p is one of the following: p is true, S has epistemically justified belief that p , and S's belief that p is undefeated. [from (4) and (5)]
7. That p is true and that S's belief that p is undefeated cannot explain why S is evidentially justified in closing inquiry. [premise]
8. Therefore, that S is epistemically justified in believing that p explains why S is evidentially justified in closing inquiry concerning whether or not p . [from (6) and (7)]
9. If that S is epistemically justified in believing that p explains why S is evidentially justified in closing inquiry concerning whether or not p , then EJ'' is true. [premise]
10. Therefore, EJ'' is true. [from (8) and (9)]

I take (2) to be highly plausible given the truth of (1). This leaves premises (1), (3), (5), (7), and (9) of the Main Argument in need of defense. In this chapter, I am only concerned with arguing for the truth of (1), or, as I referred to it earlier, TSP (The Sufficiency Premise). In chapter four, I argue in favor of (3), (5), (7), and (9) of The Main Argument. In chapter five, I articulate a theory of all-things-considered justified closure of inquiry that once again considers the justificatory status of the closure from the moral and practical domains; in other words, I consider the relationship between evidential justification and all-things-considered justification.

An Argument for TSP

There is strong intuitive evidence for TSP, and a very simple argument in its favor. The argument goes something like this. It is exceedingly strange to claim to know that p and yet affirm an intention to continue to investigate whether or not p .

The best explanation for this strangeness is that (speaking loosely) knowledge entails justified closure of inquiry.

The important point here is not merely that it seems as though one who knows that p ought to be able to quit checking whether or not p — though that's true enough. The important point is how incredibly strange it seems for one to claim to know that p while also claiming an intention to keep investigating, or the intention to keep an open mind. Kvanvig had this very point in mind when he offered the following case:

Just imagine a news conference in which the athletic director says, "We, together with the NCAA, have investigated all relevant charges of impropriety in the operation of our basketball program, and as a result, we now know that all of the charges are false and that our program is completely clean. But the investigation will continue..." Such a remark would be utterly bizarre, and the reason it would be bizarre is because knowledge involves a legitimate closure of investigation. In particular, it involves an inquiry that is of sufficient quality that it licenses the conclusion that any further learning could undermine one's present opinion only by presenting one with misleading information.³

In addition to Kvanvig's example, we might imagine a potential juror who tells the defense lawyer, "I know your client is guilty, but I'm willing to keep an open mind." I suspect that such a comment would illicit more than a few incredulous stares.

There are, of course, other explanations for the relevant strangeness. For example, one might hold that the same factors that confuse the intuitions and make infallibilism appealing are at work here, causing one to think that one's being in a position to claim to know that p entails that one should be in a position to stop investigating whether or not p . Hand-waving in the direction of the infallibilist

³ Kvanvig, "Coherentism and Justified InConsistent Beliefs: A Solution," 11.

intuition is something short of an explanation of the strangeness of claiming to know p and to intend to continue to investigate p , though it seems possible that it contains the essential ingredients for a plausible explanation. Such an explanation of the relevant strangeness, however, butts up against the intuitive evidence before us. An explanation along the lines I suggest here is an error-theory. If we can provide a theory of closure of inquiry that preserves the intuitive evidence and does theoretical work for us, that would be a good result. This strong intuitive evidence provides a cornerstone from which to begin the structure of the larger project of providing a theory that explains legitimate closure of an inquiry. If the larger project is successful, then we should be happy to be able to keep TSP; after all, it has extremely strong *prima facie* intuitive support.

To summarize, there is strong intuitive evidence that suggests that it is strange to claim to know that p and affirm an intention to continue to investigate whether or not p . The potential jurist who says that he knows that the defendant is guilty and yet insists that he is willing to keep an open mind during the trial gets strange stares. This fact is best explained by positing that knowledge entails justified closure of inquiry. Nonetheless, there are worries. Alan Millar has argued that what is needed to close inquiry is not knowledge, but reflective knowledge. Baron Reed and Jessica Brown have both offered counterexamples that might be thought to show that there are cases in which one knows and yet is not justified in closing inquiry. I'll deal with each of these objections.

Millar on Reflective Knowledge

Millar discusses closure of inquiry within the context of an argument that knowledge is more valuable than mere true belief. At first, one might be tempted to think that TSP is friendly to his project. He says, "Suppose that you are inquiring into whether something is so. Your aim is to find out whether it is so. Since finding out is nothing less than coming to know, what you aim for is knowledge."⁴ Nonetheless, Millar does not think that knowledge is sufficient for legitimately ceasing an investigation.

Millar think that one needs to not only know, but know that one knows. The need for reflective knowledge, not merely knowledge, arises for Millar because he is attempting to ground his explanation of how it is that knowledge is more valuable than mere true belief in the stability of knowledge relative to the instability of mere true belief. "Those concerned with whether something is so have good reason to want to be able to tell whether their inquiry into the matter has been successful. An inquiry is a task and agents who intend to carry out a task need to be able to tell when it has been completed."⁵

The idea seems to be that if one doesn't know that one knows what one set out to discover, then one can't reasonably stop inquiring. After all, why should one stop inquiring unless one knows that one has completed one's task? Suppose before leaving this morning I set out to discover how to get to Conway, Arkansas. Even if I come to know that the route to Conway, Arkansas is such-and-such, perhaps I

⁴ Millar, "Why Knowledge Matters," 63.

⁵ Ibid., 68.

shouldn't close inquiry into the matter until I came to know that I knew the route to Conway, Arkansas.

I once lived a summer in San Francisco. Years later, I took some friends there to visit and introduced them to a man I lived with while in San Francisco. While there, I often drove my car, following my friend, who drove his own car, to various locations in San Francisco. He would drive in such a way that it was difficult to follow him. I objected to this practice, claiming, "I don't know the way to such-and-such." He responded, "You know routes that you don't know you know." Turns out, he might have been right. He lost me on one particular occasion when I believed that I did not know where to go; but, as I drove, I remembered the way. Nonetheless, I think most people would find my objection to his habit of making it hard to follow him reasonable, since I did not know that I knew the way.

It seems to me that considerations like those of the preceding paragraph are what Millar has in mind when he writes:

The significance of reflective knowledge is that my inquiry does not merely cease with the fixation of a belief as to the location of the keys. It ceases because I know the matter is settled—the truth grasped. By contrast, with respect to mere true belief, by definition, I have no reason to think it is true and it comes with nothing else that makes sense of how the matter in hand can be settled. The same applies to a mere justified true belief, as it is, to my mind rather oddly, conceived in the framework of mainstream epistemology....Plausibly then, it is not just knowledge of the matter in hand that I am after—reflective knowledge by which I not only cease inquiry but know it to have been successfully completed because I have found out what I wanted to find out.⁶

⁶ Ibid., 68–69.

Despite the seeming plausibility of Millar's suggestion that one must have reflective knowledge before one can justifiably close inquiry, the view faces intractable difficulties. While some of these difficulties are related to the difficulties that have caused many epistemologists to give up the claim that knowledge that p entails reflective knowledge that p , Millar is quick to put out that his view does not entail that $Kp \rightarrow K(Kp)$. "It is not an implication of this view that necessarily, if you know then you know that you know. If you know that p but lack knowledge that you know that p then you will not be in a position to reap the inquiry-related benefits that can come from your knowing that p ."⁷

Nonetheless, the same problems that beset the view that $Kp \rightarrow K(Kp)$ [call this the KK thesis] beset the view that only reflective knowledge is sufficient for closure of inquiry. The worry is this: If Millar's claim were true, then a regress would make all closure of inquiry apparently unjustified.⁸ The lessons that modern epistemology has learned from the skeptic are applicable here as well. The skeptic has wielded the KK thesis as a club against all knowledge claims. If Kp entails KKp , then KKp must entail $KKKp$. There is a regress. Since it is impossible for one to comprehend, let alone know, some of the things that one must know if one can know anything, then it must be impossible for one to know anything. The best response here is just to deny the KK thesis; don't let the regress get off the ground. The first step seems to be the best step for one to deny while avoiding the charge of arbitrariness. Hence, if

⁷ Ibid., 70.

⁸ Kvanvig, "Millar on the Value of Knowledge," 93–94.

one wants to secure the possibility of knowledge, the best thing to do is to deny the KK thesis.

Once one becomes aware of the regress problem, it is easy to see how a similar regress makes Millar's condition on legitimate closure of inquiry seemingly untenable. After all, by Millar's own lights, your knowledge that you know that p can't reap the inquiry related benefits that come from your knowing that you have knowledge that you know that p . It seems, then, that if Millar's condition on legitimate closure of inquiry were true, there would be no legitimate closure of inquiry.

Millar has a reply: He might claim that one is not inquiring into whether or not one has reflective knowledge. Millar's claim is this: one must have reflective knowledge before one is justified in closing inquiry. The skeptic might ask Millar, "How does one know that one has reflective knowledge?" Millar might reply, "Suppose that one doesn't know that one has reflective knowledge. So? As long as he isn't inquiring into whether or not he is justified in ceasing inquiry, that's okay. I never said that he couldn't have reflective knowledge without having third order knowledge."

While this escape route might give him some room to work, there nonetheless seems to be tension here between two claims that Millar needs to endorse to avoid the charge of vicious regress. The tension is between the following two claims: a) one is never justified in closing inquiry until one knows that one has completed one's task, and b) one can be justified in closing inquiry even when one doesn't know that one knows that one completed one's task. Given that Millar

thinks that only reflective knowledge provides one with justification for easing the irritation of one's doubt, why does he think that one has reason to be free of doubt regarding one's reflective knowledge? Perhaps this tension does not conclusively demonstrate the case against reflective knowledge as a condition for justified closure of inquiry.

Millar has another problem. His view also returns the verdict that many closures of inquiry are illegitimate when strong intuitive evidence suggests they are legitimate. When I got up this morning, I inquired into where I left my shoes last night. For a moment, I couldn't remember. Then, I remembered that I left them in the living room. I quit inquiring into where my shoes were, and I did so without ever believing that I knew where my shoes were; it never occurred to me to believe that I knew where my shoes were. But if I did not believe that I knew where my shoes were, then I could not have known that I knew where my shoes were. If Millar's thesis is true, I was unjustified in closing inquiry. Surely, however, I was justified in closing inquiry.

Millar might weaken his requirement, claiming that one need not actually know that one knows, but only be disposed to believe (in the right kind of way) that one knows, were it to occur to one to consider the question. But even this is too strong. After I put my shoes on (they were in the living room after all), my three-year old daughter asked me where her doll, Clarise, was. I told her to look around. She found Clarise in her room, and immediately quit looking to see where her doll was. I assume that she did so legitimately, and she did so without even having

the conceptual resources to consider whether or not she knew that she knew where her doll was.

Potential Counterexamples TSP

TSP is similar to a principle discussed by Jeremy Fantl and Matthew McGrath (hereafter: F&M). In *Knowledge in an Uncertain World*, they defend

KJ If you know that p , then p is warranted enough to justify you in φ -ing, for any φ .

In defending KJ, they take themselves to also be defending:

Action If you know that p , then if the question of whether p is relevant to the question of what to do, then it is proper for you to act on p .

Best Results If you know that A will have the best results of your available options, then you are rational to do A.

Preference If you know that p , then you are rational to prefer as if p .

Inquiry If you know that p , then you are proper not to inquire further into whether p .

One might immediately notice the similarity between (Inquiry) and TSP.

TSP If S knows p , then S is evidentially justified in closing inquiry concerning whether or not p .

The difference between TSP and Inquiry is that TSP plainly qualifies the sense in which one is justified in closing inquiry. The qualification (if, in fact, there is any similar qualification) is less obvious in Inquiry.

Since I have stipulated that closing inquiry is an action, the truth of KJ entails the truth of TSP, given that the sense of justification in KJ entails that one is evidentially justified. However, the truth of TSP does not entail the truth of KJ.

F&M's claim is stronger than my own.

Nonetheless, some of the objections to KJ also serve as potential objections to TSP. I consider two of them here. Barron Reed has raised one such objection.

Reed presents us with the following potential counterexample:

Punishment/reward case 1. You are participating in a psychological study intended to measure the effect of stress on memory. The researcher asks you questions about Roman history—a subject with which you are well acquainted. For every correct answer you give, the researcher will reward you with a jelly bean; for every incorrect answer, you are punished by an extremely painful electric shock. There is neither reward nor punishment for failing to give an answer. The first question is: when was Julius Caesar born? You are confident, though not absolutely certain, that the answer is 100 BC. You also know that, given that Caesar was born in 100 BC, the best thing to do is to provide this answer....⁹

Reed articulates the above case as a counterexample to an earlier, less general claim articulated by F&M. F&M claim that if you know that *p*, then you should be able to act as if *p*.¹⁰ Reed, however, takes the above case to be a counterexample to this claim. Even if you know that Julius Caesar was born in 100 BC, you ought not answer the question, since the reward is so small, and the price of being wrong so high.

F&M seem to claim that cases like the above simply take a small, usually idle, chance of error and manipulate the stakes such that the idle chance of error becomes salient. If the otherwise idle chance of error becomes salient, then knowledge is lost. Some people think it is a bad feature of F&M's view that knowledge can be gained or lost merely from practical consideration. Reed goes on to offer another a second case

⁹ Reed, "A Defense of Stable Invariantism," 228–229.

¹⁰ Fantl and McGrath, "Evidence, Pragmatics, and Justification," 67.

Punishment/reward case 2. As before, a researcher is asking you questions about Roman history. But now there are two punishment/reward scenarios. The first is as before: for a correct answer, you receive a jelly bean, and for an incorrect response, you get a very painful electric shock. In the second, a correct answer results in a \$1000 reward, while an incorrect one leads only to a very mild slap on the wrist. In both scenarios, there is neither reward nor punishment for failing to answer. Although you must consider both scenarios simultaneously, you are not bound to give the same answer in each.¹¹

In the second case, one seems rational in answering the question. In the first case, one does not seem rational in answering the question. In both cases, Reed says, the agent knows the answer.

Under F&M's view, it is part of the nature of knowledge that it is sensitive to practical considerations in ways that make it impossible for practical rationality and knowledge to come apart in the way that the first case suggests. If they are right, then it seems that either you do not know that Julius Caesar was born in 100 BC in the first case (despite what the case says), or that it is rational (contra the putative intuitive evidence) for you to answer the question.

Notice that nothing in KJ or any of the more specific principles entails a theory of rationality in which rationality obligates one to answer the question, even given that one knows the answer and one prefers having the jelly bean to not having the jelly bean. It seems theoretically open to a proponent of KJ to soften the blow of

¹¹ Reed, "A Defense of Stable Invariantism," 230.

Reed's counterexample by claiming that rationality permits, but does not require, one to answer the question.¹²

Notice also that the defender of KJ might utilize the distinction between alethic and epistemic justification.¹³ Suppose that the person in Reed's first case does not know that Julius Caesar was born in 100 BC (despite what the case claims). Suppose it merely seems right to say that such a person does know. Perhaps our intuitions are clouded by our inability to distinguish between alethic and epistemic justification. Both the subject in the first case and the subject in the second case have a justified true belief that Julius Caesar was born in 100 BC. Because of this similarity, we are tempted to think they both know that Julius Caesar was born in 100 BC. Suppose, though, that only the person in the second case is justified in believing that she knows that Julius Caesar was born in 100 BC. In that case, then (supposing Kvanvig's distinction between alethic and epistemic justification is true) only the person in the second case has epistemic justification for the claim in question, and hence only she knows that Julius Caesar was both in 100 BC.

While I find the above response to have merit, it is more than I need to defend. I need not make so strong a claim as KJ in order to defend TSP. Even given that the subject of Reed's first case a) does know that Julius Caesar was born in 100 BC and b) is not rational in answering the question, it does not follow that TSP is

¹² I am not sure that F&M would be happy with this response, since they understand justification as imposing an obligation, where I am working with a permissive conception of justification.

¹³ This reply is not open to F&M, as they favor The Equivalence Thesis: *p* is knowledge-level justified for you iff you are justified in believing that *p*. Fantl and McGrath, *Knowledge in an Uncertain World*, 98.

false. Reed's case is not a case of closing inquiry. The question at issue is whether or not the person in the case is rational in answering the question, not in whether or not the person is rational in closing inquiry. While answering a question, like closing inquiry, is an action, it is not the same type of action. It is possible that legitimately closing an investigation does not operate under the same normative constraints as does the action of rationally answering a question. In order to effectively challenge TSP, one needs a case in which it appears that one knows p , but is not justified in closing inquiry concerning whether or not p . Jessica Brown has given such a case. I turn to it now.

Jessica Brown and the Diligent Surgeon

Jessica Brown has argued that knowledge that p is not always sufficient to close inquiry concerning p . She gives the following case:

The Surgeon Case: A student is spending the day shadowing a surgeon. In the morning he observes her in clinic examining patient A who has a diseased left kidney. The decision is taken to remove it that afternoon. Later, the student observes the surgeon in theatre where patient A is laying anaesthetized on the operating table. The operation hasn't started as the surgeon is consulting the patient's notes. The student is puzzled and asks one of the nurses what's going on:

Student: I don't understand. Why is she looking at the patient's records? She was in clinic with the patient this morning. Doesn't she even know which kidney it is?

Nurse: Of course, she knows which kidney it is. But, imagine what it would be like if she removed the wrong kidney. She couldn't operate before checking the records.¹⁴

¹⁴ Brown, "Subject-Sensitive Invariantism and the Knowledge Norm for Practical Reasoning," 176.

Brown takes the above story to support the intuition that, as the nurse says, the surgeon does know that the left kidney is the diseased kidney; nonetheless, the surgeon seems unjustified in closing inquiry in the matter. This case, unlike Reed's jelly bean case, puts direct pressure on TSP. If TSP is true, then if the doctor knows which kidney is diseased, then the doctor is evidentially justified in closing inquiring regarding whether or not *p*.

F&M think that Reed's and Brown's examples "(at best) reveal that the principles which they are meant to refute are not clearly true, prior to theorizing...."¹⁵ Their reason for holding that these counterexamples are indecisive is that the examples can be modified, without changing the subjects' strength of epistemic position, to "pump intuitions for the opposite conclusion."¹⁶ For example, they modify the surgeon case as follows:

Student: I don't understand. Why is she looking at the patient's records? She was in clinic with the patient this morning. Doesn't she even know which kidney it is?

Nurse: She did check the patient this morning, but every good surgeon checks the patients' chart so that she *knows* which kidney to take out.

It seems to me certainly correct that the above conversation seems just as plausible (plausible, that is, not merely in terms of the way we talk, but also in terms of its proper use of 'knows') as Brown's. F&M seem to insist that the surgeon does not know which kidney is diseased, despite the nurse's claim that the surgeon does know. For this response to work, F&M need an error theory. The reason that the nurse's attribution of knowledge regarding the surgeon seems credible is because

¹⁵ Fantl and McGrath, *Knowledge in an Uncertain World*, 63.

¹⁶ Ibid., 62.

many people find it intuitive that the surgeon does know. How might one explain this error in intuition?

Here is a potential error theory: many people associate justified (even when justification is merely alethic) true belief with knowledge. Since people see that the surgeon has a justified true belief regarding which kidney is diseased, they come to believe that the doctor knows which kidney is diseased. However, one might think, if one reflected more on the inadequacy of mere alethic justification to produce knowledge, one might be less inclined to error in this way.

Alternatively, it might not be people's intuition that the surgeon knows that is in error; it might be their intuition that she should check. In this regard, people's intuition might be led astray by the human tendency to availability bias. The availability bias is a phenomenon by which people tend to exaggerate the danger associated with risks that are more available to them (for example, in virtue of being covered by the media more often, or more easily seen, etc) and to minimize the danger associated with risks that are less available to them. In the surgeon case, most of us quickly realize the danger of removing the wrong kidney; we've heard of those kinds of cases. The danger associated with taking *too much* care is unseen. We don't worry about a hospital or a doctor that spends *too many* resources on being sure that a mistake is not made. But perhaps we should. A doctor who is able to help fewer people, or who takes so long helping people that she becomes mentally fatigued, is a real danger – though not a particularly available one (in the sense explained above). Hence, it is possible that people who hear Brown's surgeon case are in error when they say that she still needs to check; unsurprisingly, given

what we know about the availability bias, they are spending too much time worrying about the risk associated with her taking too little care. If they would spend more time worrying about the risk associated with her taking too much care, they would not be inclined to think that she should check.

Here is another reason to think that the nurse is mistaken in her claim that the doctor should check. Imagine the following conversations:

Student: I don't understand. She just checked the records. Why is she checking them again? Doesn't she know which kidney it is?

Nurse: Of course, she knows which kidney it is. But, imagine what it would be like if she removed the wrong kidney. She couldn't operate before double-checking the records.

Or, perhaps the following:

Student: I don't understand. She just double-checked the records. Why is she checking them again? Doesn't she know which kidney it is?

Nurse: Of course, she knows which kidney it is. But, imagine what it would be like if she removed the wrong kidney. She couldn't operate before triple-checking the records.

Obviously, there is some point at which the Nurse's defense of the surgeon's checking the records becomes implausible. It seems as though the least arbitrary point at which the surgeon is justified in closing inquiry is the point when the surgeon comes to know which kidney is diseased.

A second reply that F&M can make to the Brown case is to concede that the surgeon does know which kidney is diseased, and that she should check the records, but to insist that the surgeon's reasons for checking the records are unrelated to the strength of her epistemic position. F&M say that if one knows that *p*, then *p* is *warranted enough* to justify one in φ -ing. They go on to say that "*p* is warranted

enough to justify you” means “no weaknesses in your epistemic position with respect to p—in your position along any of the truth-relevant dimensions—stand in the way of p justifying you.”¹⁷ F&M might hold that doctors have an ethical responsibility to double-check in situations like the one in Brown’s case; many hospitals might well require it as a matter of procedure. This would give her an ethical responsibility to check the records, even given that she knows which kidney is diseased. Given F&M’s understanding of ‘warranted enough’, such a fact would not count against KJ, because (in such a case) it is not a weakness in the surgeon’s epistemic position that makes her unjustified in failing to check the records.

It is possible, of course, to modify the example in such a way that removes the role responsibility that the surgeon (given her role as a surgeon) has to check the records. Allison Thorton has suggested such a case:

Anaphylactic Shock: Annie and her sister Susan are spending the evening watching reruns of *The Bachelor*. They have purchased a plethora of junk food to get them through the episodes, including a bag of plain M&Ms. Unbeknownst to either of them, there is a rogue peanut M&M mixed in with plain ones. This is a problem because Susan has a life-threatening allergy to legumes. Unfortunately, it is Susan who picks the peanut M&M from the mix. She chews it, noticing a funny taste, but she doesn't realize the danger until she swallows, at which point her body goes into anaphylactic shock. She recognizes the symptoms and administers medication to slow the reaction. But that is a temporary _x. Annie needs to get her to the hospital right away. They hurry to her car. Annie knows how to get to the hospital. In fact, she was there earlier this morning delivering flowers to a recovering friend. But she doesn't have time to take a wrong turn or miss an exit. She has ten minutes before Susan's airway swells shut completely and thirteen minutes until it's too late. The hospital is at least nine minutes away. As Annie and Susan rush to the car, Annie takes out

¹⁷ Ibid., 66.

her iPhone and tells Siri to provide directions to the hospital. Confused, Susan asks what's going on:

Susan: I don't understand. Why are you getting directions? Weren't you there this morning? Don't you even know how to get to the hospital?

Annie: Of course I know how to get to the hospital. But imagine what would happen if I took a wrong turn or missed an exit!¹⁸

Thorton's case is a case of procedural knowledge, or knowledge how. One might worry that procedural knowledge does not reduce to propositional knowledge. Nonetheless, it seems simple enough to modify the example so that it deals straightforwardly with propositional knowledge. We might imagine Annie calling a friend and asking her friend if a certain highway (say, highway 6) is the way to the hospital. She knows, we might suppose, that highway 6 is the way to the hospital, and yet she is not justified in closing inquiry concerning <Highway 6 is the way to the hospital.>

Thorton's case is less obvious than the surgeon case – perhaps because one's responsibilities *qua* friend are less obvious than one's responsibilities *qua* surgeon. Further, it seems possible that Annie's responsibilities as a friend morally require her to double check, even if her evidence justifies her in closing the investigation.

Thorton recognizes this:

It is important to note that there may still be a policy about double-checking in play in this case. The difference between this policy and the objectionable policy present in Brown's example is that the former is imposed by someone on herself and the latter by an institution on individuals over whom it has some authority. The institution can create for those individuals ulterior motivation for double-checking that is unrelated to epistemic position. It's unlikely that an individual, or at least an individual in Annie's situation,

¹⁸ Thorton, "Knowing and Double Checking."

creates such epistemically irrelevant motivation for double-checking. Rather, it seems that Annie's reason for double-checking seems to be directed at a weakness in her epistemic position with respect to how to get to the hospital.

It is not clear to me why we should think that the responsibilities one places on oneself are relevant to the strength of one's epistemic position, but the responsibilities imposed on one by someone else is unrelated to epistemic position.

Regardless, I need not defend KJ. The truth of TSP can withstand the triadic conclusion that Annie knows that the highway is the way to the hospital, that Annie is not justified in failing to check to see if the highway is the way to the hospital, and that Annie's lack of justification in closing inquiry is due to a weakness in her epistemic position. Such a conclusion would damn KJ, but not TSP. This is because TSP only offers a sufficiency condition for an *evidentially* justified closure of inquiry, not an *all-things-considered* justified closure of inquiry; it does not entail that knowledge is sufficient for all-things-considered justification. It might be that what one is evidentially justified in doing can come apart from what one is practically or morally justified in doing. In chapter four, I will consider the manner in which all-things-considered justification to close inquiry is sensitive to cases in which one's practical interest comes apart from what one evidence justifies.

Epistemic and Alethic Justification: A Helpful Digression

Recall that I am defending TSP as a first step in the defense of EJ''.

EJ'' S's evidence justifies S in closing inquiry concerning whether or not *p*
 iff S's belief that *p* (or S's belief that $\sim p$) is epistemically justified.

Before continuing in my defense of EJ'', I need to provide myself with a resource that handles one obvious objection. The resource I need is a distinction between at least

two different kinds of truth-oriented justification. Without a distinction between epistemic justification and at least one other kind of truth-oriented justification, there is an intractable difficulty for EJ". In order to illustrate this problem, consider that you hold in your hands a lottery ticket. The ticket in your hand has a 1:1,000,000 chance of winning for you something that you value greatly. Here are two extremely plausible claims about this situation: first, you are justified in believing that your ticket is not the winning ticket and, second, your evidence does not justify closing inquiry into whether or not your ticket is the winner. The sense in which you are justified in believing that your ticket is not the winner is clearly truth related. Thus, for EJ" to be plausible, I need to hold that there is a kind of justification that, while oriented to the truth, is not epistemic justification.

Jonathan Kvanvig has argued for just this conclusion.¹⁹ He claims that the kind of justification that, when it is present with undefeated true belief, yields knowledge is specifically epistemic justification, and that a second type of common, truth-oriented justification can attach to undefeated true belief without yielding knowledge; the latter justification he calls alethic justification.

The chief reason for thinking that alethic justification exists is that there are cases in which, intuitively, one has undefeated justified true belief that *p* and yet does not know that *p*. The lottery example is an excellent illustration of such a case. In the lottery example, one is justified in believing that one's ticket will lose, and yet, even in cases when one's belief that one's ticket will lose is undefeated and true, one

¹⁹ Kvanvig, "Coherentism and Justified Inconsistent Beliefs: A Solution," 30–31.

does not know that one's ticket will lose. The claim that one does not know that one's ticket will lose is supported in two ways.

First, we note that if one is justified in believing about a single ticket that it is not the winner, then one is justified (in the right conditions) in believing of each ticket individually that it is not the winner. But, then, it would seem that one should be justified in believing the conjunction of all such claims. But one is certainly not justified in believing that all of the tickets will lose. This consideration is not conclusive, since it is open to one to hold that one can know the conjuncts without knowing the conjunction; the conjunction has a decreasing chance of being true as additional individual conjuncts (ie, *this* ticket a loser, and *this* ticket is a loser) are added to it. At some point, one might hold, probability is degraded to the point that knowledge is lost.

A second consideration that supports the claim that one does not know that one's ticket is a loser is the fact that one would be irrational if one gave away one's ticket before the winner was announced. But if you knew that your ticket was a loser, you would not be irrational in giving your ticket away. So, you must not know that your ticket is a loser. Since you do have a justified true belief that your ticket is a loser, and the relevant justification is a truth-oriented justification, then (given the traditional view of knowledge), there must be another kind of truth-oriented justification that is not epistemic justification.

The distinction between alethic and epistemic justification, if it exists, explains how it is that one can have undefeated, justified true belief that one's lottery ticket is a loser without knowing that one's lottery ticket is a loser. The

distinction itself is not what makes TSP plausible, however. What makes it plausible is primarily the strong intuitive evidence itself. It just seems that when one knows something, one ought to be able to quit checking to see if it is true. Even stronger, it seems exceedingly strange for one to claim to know something and yet keep checking to see if it is true.

By my lights, the intuitive evidence that suggests that knowledge entails justified closure of inquiry is strong. If TSP is true (as the above intuitive evidence suggests), then it is easy to see why one might think there is a distinction between epistemic and alethic justification. In the lottery case, it is clear that one is justified (in some sense) in believing that one's ticket lost the lottery. But, if TSP is true, one does not know that one's ticket will lose the lottery – even in cases in which one's ticket is a loser.

Kvanvig goes on to give the following (rough, later revised)²⁰ distinction between epistemic and alethic justification:

If you justifiably believe that p where the kind of justification in question is epistemic justification, then you have justification for believing that you know that p . If you justifiably believe that p , but lack justification for believing that you know, then you have ordinary justification only.²¹

Here, when Kvanvig speaks of ordinary justification, he means what he elsewhere

²⁰ The later revision was chiefly aimed at concerns that one can be mistaken about what one believes. In his revised distinction, he doesn't quite claim that justification for believing that one knows p follows from one's having epistemic justification for p . Instead, he claims that justification (of the propositional sort) for believing that the non-psychological conditions for knowledge hold in regards to p follows from one's having epistemic justification for p . For my purposes, I'll ignore this subtlety.

²¹ Kvanvig, "Coherentism and Justified InConsistent Beliefs: A Solution," 31.

refers to as alethic justification. As a result, all references to justification in the above formulations that are not specified as of the epistemic variety should be understood as of the alethic variety only.

There is a problem with Kvanvig's articulation of the distinction between alethic and epistemic justification. The problem is this: some smart, informed people believe that they know that their lottery ticket will lose, despite the fact that they admit that they ought not give the ticket away. I don't want to claim that such people are unjustified in their belief that they know that their ticket will lose; I think they are wrong, not irrational. People who hold such a view do so not primarily because of the intuitive case in favor of the claim that people know their lottery ticket is a loser²², but because accepting this result provides them with theoretical consistency and seems consistent with fallibilism, which itself is needed in order to avoid skepticism. I recognize that such people have reasons for believing that they know their ticket is a loser. I just happen to think that the reasons for believing that we don't know that our lottery ticket is a loser are better reasons; namely, one can still achieve theoretical consistency without sacrificing highly intuitive principles. Despite the difficulties in specifying precisely the distinction between epistemic and alethic justification, it seems plausible that some such distinction exists.

Conclusion

In this chapter and the next, I am arguing in favor of

EJ'' S's evidence justifies S in closing inquiry concerning whether or not p
iff S's belief that p (or $\sim p$) is epistemically justified.

²² This is not to say that nobody claims to find it intuitive that people know such things. Some people might find that intuitive; I certainly do not.

My Main Argument, the conclusion of which is EJ'', is as follows:

1. If S knows *p*, then S is evidentially justified in closing inquiry concerning whether or not *p*. [premise]
2. Either S's knowledge that *p* or a fundamental constituent of S's knowledge that *p* explains why S is evidentially justified in closing inquiry concerning whether or not *p*. [premise]
3. S's knowledge that *p* can't explain why S is evidentially justified in closing inquiry concerning whether or not *p*. [premise]
4. A fundamental constituent of S's knowledge that *p* explains why S is evidentially justified in closing inquiry concerning whether or not *p*. [from (2) and (3)]
5. The fundamental constituents of S's knowledge that *p* are the following: *p* is true, S has epistemically justified belief that *p*, and S's belief that *p* is undefeated. [premise]
6. The feature of S's knowledge of *p* that explains why S is evidentially justified in closing inquiry concerning *p* is one of the following: *p* is true, S has epistemically justified belief that *p*, and S's belief that *p* is undefeated. [from (4) and (5)]
7. That *p* is true and that S's belief that *p* is undefeated cannot explain why S is evidentially justified in closing inquiry. [premise]
8. Therefore, that S is epistemically justified in believing that *p* explains why S is evidentially justified in closing inquiry concerning whether or not *p*. [from (6) and (7)]
9. If that S is epistemically justified in believing that *p* explains why S is evidentially justified in closing inquiry concerning whether or not *p*, then EJ'' is true. [premise]
10. Therefore, EJ'' is true. [from (8) and (9)]

In this chapter, I argued for the truth of (1), which I called TSP. I showed how intuitive evidence gives us good reason to accept TSP. Namely, it is exceedingly strange to claim to know that *p* while also claiming that you will keep an open mind about *p*. After providing a brief argument for TSP, I considered objections. First, I considered Millar's claim that reflective knowledge is necessary for closing inquiry. Second, I considered putative counterexamples to TSP. I showed how even claims stronger than TSP, like F&M's KJ thesis, have responses to such supposed counterexamples. Then, I showed how, even given that the counterexamples defeat KJ, TSP makes no claim about all-things-considered justification. Given the strong

intuitive support for TSP and the responses available to the counterexamples that have been discussed in the literature, one should accept the principle. Finally, I provided resources for the road ahead by distinguishing between alethic and epistemic justification. In this next chapter, I will argue for the truth of (3), (5), (7), and (9) of the main argument

CHAPTER FOUR

Closure of Inquiry and the Sufficiency of Epistemic Justification

In chapter two, I argued that justified closure of inquiry requires a role for epistemic normativity. In chapter three, I began an argument in favor of a claim that, if true, provides some clarity regarding the relationship between justified closure of inquiry and epistemic normativity. Recall EJ'':

EJ'' S's evidence justifies S in closing inquiry concerning whether or not p
 iff S's belief that p (or S's belief that $\sim p$) is epistemically justified.

I am attempting to demonstrate the following, which I have called the Main

Argument:

1. If S knows p , then S is evidentially justified in closing inquiry concerning whether or not p . [premise]
2. Either S's knowledge that p or a fundamental constituent of S's knowledge that p explains why S is evidentially justified in closing inquiry concerning whether or not p . [premise]
3. S's knowledge that p can't explain why S is evidentially justified in closing inquiry concerning whether or not p . [premise]
4. A fundamental constituent of S's knowledge that p explains why S is evidentially justified in closing inquiry concerning whether or not p . [from (2) and (3)]
5. The fundamental constituents of S's knowledge that p are the following: p is true, S has epistemically justified belief that p , and S's belief that p is undefeated. [premise]
6. The feature of S's knowledge of p that explains why S is evidentially justified in closing inquiry concerning p is one of the following: p is true, S has epistemically justified belief that p , and S's belief that p is undefeated. [from (4) and (5)]
7. That p is true and that S's belief that p is undefeated cannot explain why S is evidentially justified in closing inquiry. [premise]
8. Therefore, that S is epistemically justified in believing that p explains why S is evidentially justified in closing inquiry concerning whether or not p . [from (6) and (7)]

9. If that S is epistemically justified in believing that p explains why S is evidentially justified in closing inquiry concerning whether or not p , then EJ'' is true. [premise]
10. Therefore, EJ'' is true. [from (8) and (9)]

In chapter two, I argued that TSP (1) is true. I still need to argue in support of (2), (3), (5), (7), and (9). In this chapter, I take up that task. After completing my defense of the Main Argument, I give another upshot of EJ''; it does much of the theoretical work of a similar principle defended by Fantl and McGrath, without falling prey to a compelling counterexample.

Premise 2

A defense of (2) can be handled quickly. Since we know that knowledge entails evidentially justified closure of inquiry, we know that one of the fundamental constituents of knowledge, or some subset of the fundamental constituents of knowledge (perhaps the subset including all of the these fundamental constituents; ie, knowledge itself), entails evidentially justified closure of inquiry. I offer no analysis of what it means for one thing to explain another; I trust that a pre-theoretical understanding of explanation will suffice.

Premise 3

If knowledge itself is what evidentially justifies a closure of inquiry (and not merely some proper subset of the fundamental constituents of knowledge), then one should expect to find cases in which removing any of the fundamental constituents of knowledge, while holding everything else fixed, would cause the relevant agent to lose her evidential justification to close inquiry. But when we subtract any one of

the fundamental constituents of knowledge (save epistemic justification), we find just the opposite.¹

Fantl and McGrath make a similar point in what they call the subtraction argument.² They give a case in which our car has broken down on a bitterly cold night, in the middle of nowhere. We need to find shelter. After walking for hours, we clearly see a barn about 50 yards to our right.

In such a case, we are justified in closing our inquiry. We would walk to the barn on our right, but we would not likely walk to the barn with the intention of inquiring into “there’s a barn.” We’d describe our action as “going to the barn to get shelter” not as “inquiring into whether there is a barn.” After all, we clearly see the barn. We know that a barn is there. No more inquiry is needed.

If we vary the case such that we have been Gettiered (the barn that we see is really just an extremely realistic-appearing barn façade right in front of an actual barn, or some such story intended to invoke the Gettier mechanism), there is no change in the status of the evidential justification of our closure of inquiry. We were evidentially justified in closing inquiry in the Gettier case, and we remain evidentially justified in closing inquiry in the normal case. Moreover, if we vary the

¹ Since a successful defense of TSP does not establish that knowledge is a necessary condition for one’s closure of inquiry to be evidentially justified, it is not quite enough for me to merely show that there are cases in which we can subtract one of the fundamental constituents of knowledge without causing the closure to lose its status as justified. Nonetheless, if I can show that the connection between the various fundamental constituents of knowledge (with the exception of epistemic justification) and evidentially justified closure of inquiry is such that the removal of their status never seems to effect the justificatory status of the closure, then I have presented strong evidence against the claim that knowledge itself is what evidentially justifies closure of inquiry.

² Fantl and McGrath, *Knowledge in an Uncertain World*, 99.

case such that there is no barn to the right (only extremely realistic-looking barn façade), there is still no change as to the status of the evidential justification of the closure of inquiry in question. External defeat (or whatever is the missing ingredient in the Gettier cases) and/or truth are not the right kinds of things to cause one to lose one's evidential justification to close inquiry.

Premise 5: The Easier Objections

The claim that the fundamental constituents of knowledge are epistemically justified belief, truth, and the appropriate Gettier condition is controversial. With the exception of the truth condition, each condition has been challenged. For example, Colin Radford argued that knowledge does not entail belief.³ Crispin Sartwell has argued that knowledge is merely true belief, with no fundamental connection to justification or the Gettier condition.⁴ Such views, while present, are firmly in the minority. It is easy to see why.

Regarding Radford's claim, the proposed cases of knowledge in the absence of belief can be adequately explained as not cases of knowledge. The kinds of cases Radford has in mind are cases in which one seems to know a proposition (because one answers it correctly on a test), but does not believe the proposition.⁵ One of his examples involves someone named Jean taking an English history quiz. Jean is asked when Elizabeth died. She has just missed a number of similar questions, and

³ Radford, "Knowledge—by Examples."

⁴ Sartwell, "Knowledge Is Merely True Belief."

⁵ Radford, "Knowledge—by Examples," 2.

does not have any belief regarding the specific year Elizabeth died. Nonetheless, she answers “1603” and gets the question correct. There is a further implication that Jean was vaguely remembering something she had heard before. Radford’s claim is that Jean knows that Elizabeth died in 1603 even though Jean did not believe that Elizabeth died in 1603.

Most epistemologists reject the claim that Jean knows, in the relevant sense, that Elizabeth died in 1603. Keith Lehrer, for example, concedes that there is a sense in which the person who gets the test question right without believing the answer knows the answer (after all, it was not just luck that caused him to choose the right answer; he remembered it), but argues the sense in which he knows the correct answer is not the sense of robust knowledge that concerns the epistemologist.⁶ The student *once* knew the correct answer in the robust sense of ‘knew’, but that knowledge has vanished; in its place, there remains merely a cognitive connection that explains how the student chose the correct answer. We are no more compelled to explain *how* the student knows the correct answer by saying “because he knew it” than we should feel compelled to answer “because he knew it” in cases when the student got the *wrong* answer. The causal cognitive connection itself is not enough demonstrate the presence of knowledge in the robust sense, even in cases in which the student happened to get it right.

As for Sartwell’s claim that knowledge is merely true belief, with no fundamental connection to justification or the Gettier condition, we would do well to notice that there seem to be many unjustified true beliefs that fail to count as

⁶ Lehrer, *Theory Of Knowledge*, 37.

knowledge. It is important to note that Sartwell's point doesn't have to do with the externalism/internalism debate. He adds no condition about reliability or proper-functioning faculties. He simply thinks that all true beliefs, regardless of any external factors, count as knowledge. But this seems false. If I were to believe, truly, that there is an even number of stars in the universe, I would still not know it.

While few people think true belief is sufficient for knowledge, there are two more common reasons for rejecting the claim that epistemic justification is a constituent of knowledge. The first and chief concern is the possibility that some external relation must hold (perhaps about the reliability of the belief-producing mechanism, or the proper-functioning of the faculties that led to the belief) for true belief to result in knowledge, and that justification is not an additional constituent beyond that external relation. Responding to such worries will be the largest task of this section. Before I respond to this worry, though, I want to explain a second reason for rejecting the claim that epistemic justification is a constituent of knowledge

The second reason has to do with the distinction between alethic and epistemic justification. Recall (5):

5. The fundamental constituents of S's knowledge that *p* are the following: *p* is true, S has epistemically justified belief that *p*, and S's belief that *p* is undefeated.

(5) entails that a proposition, in order to be known, must be epistemically justified. Recall also that my understanding of epistemic justification implies a distinction between epistemic justification and mere alethic justification. Hence, one might reject (5) because one thinks that justification sufficient for belief (mere

alethic justification) is sufficient for knowledge; stronger justification is not needed.

I will handle this worry first. Then, I will handle the worry about externalism.

The Main Argument can be restated under the assumption that alethic justification is sufficient for knowledge, and it would still go through. The salient issue here is the truth of what Fantl and McGrath (hereafter: F&M) call the Equivalence Thesis.⁷

The Equivalence Thesis p is knowledge-level justified for you iff you are justified in believing that p .

If, as I argued in chapter two, there is a distinction between alethic and epistemic justification, then The Equivalence Thesis is false. Nonetheless, many people *seem* to accept it; that is to say, the idea that $K=JTB+$ *seems* to presuppose that The Equivalence Thesis is true, since the proponents of $K=JTB+$ don't often speak in terms of two varieties of justification.

By my lights, the truth of TSP makes The Equivalence Thesis almost certainly false. To see this, notice that the truth of TSP (or KJ) entails that I am justified in not checking to see if my lottery ticket won in all cases in which I have knowledge-level justification for believing that my lottery ticket did not win. It seems very likely that I am justified in believing that my lottery ticket did not win, but far less likely that I am justified in closing inquiry on the matter. To put the matter another way, it is easier to hold that I am justified in believing, though I do not know, that my lottery ticket is not a winner than it is to hold that I am not justified in believing that my lottery ticket is not a winner. Given the incredibly low odds that my ticket is a

⁷ Fantl and McGrath, *Knowledge in an Uncertain World*, 98.

winner, certainly I am justified in believing my ticket is a loser—though I don't know it (as the truth of TSP suggests). In most cases, I couldn't withhold belief on that one even if I tried.

Nonetheless, there is pressure to accept The Equivalence Thesis. After all, the JTB+ account of knowledge does seem to presuppose it, or at least assume it. Suppose The Equivalence Thesis is true. Then, premise 5 (*as I mean it; wherein there is an implied distinction between epistemic and alethic justification*) is false, or at least problematically confusing, because there is no true distinction between alethic and epistemic justification. Okay; just modify the argument so that all references to epistemic justification are changed to whatever kind of justification that you think is constitutive of knowledge. The argument still goes through. As such, we need not be concerned that (5) is false when it is taken to imply that mere alethic justification is insufficient for closing inquiry. We can simply modify the argument.

Premise 5: The Radical Externalist Worry

The bigger worry is the claim that knowledge is fundamentally constituted by truth, belief, and some proper external relationship. Many philosophers (following Plantinga) refer to this external relationship as warrant. While warrant must be added to true belief to get knowledge, it is not always a reason for the agent to believe the claim in question. For example, some philosophers think knowledge

is reliably formed true belief⁸, others that it is true belief that is the result of proper functioning faculties operating in the environment for which they were intended.⁹

Such views, views that insist that there are non-internal conditions on knowledge beyond just the truth condition and the Gettier condition, are broadly understood as externalist theories of knowledge. On such views, whether one knows a proposition is a function of the fact that one has a true belief that stands in a certain external relation to the world. Under many such views, knowledge does not require justification; it only requires that the relevant external relations hold, and that the agent believe the true proposition in question. Given the truth of The Sufficiency Premise, such views are implausible.¹⁰

The reason that the claim that knowledge does not require justification (following F&M, I will call such views radical externalism) seems incompatible with the truth of TSP is that the truth of radical externalism seems to entail that you can know a proposition while being justified in believing that it is false. This is because, on most external theories of knowledge, it seems as though the relevant external objective relation to the world could hold, and an agent could believe the true proposition in question, despite the fact that one's epistemic situation is such that one is propositionally justified in believing the denial of the proposition. First, I will support my claim that radical externalism entails that one can know a proposition

⁸ Alvin Goldman, "What Is Justified Belief?".

⁹ Plantinga, *Warrant and Proper Function*.

¹⁰ F&M make use of a very similar argument in defending their own views from externalist theories. See chapter four of *Knowledge in an Uncertain World*.

while being justified in believing that the proposition is false. Then, I will show why this fact makes radical externalism incompatible with TSP.

Laurence Bonjour introduced the case of Norman and the crystal ball.¹¹ In this case, Norman looks at a crystal ball and comes to believe that the president is in New York. Norman has no better reason to think that the crystal ball is a reliable guide to truth than you or I do, and he has no other reason to suppose that the president is in New York. Nonetheless, it turns out to be the case that the crystal ball is reliable. According to some versions of radical externalism, Norman knows that the president is in New York. Since we have seen that TSP is true, it follows that Norman is evidentially justified in closing inquiry concerning whether or not the president is in New York. But, in this case, Norman has very little evidence that the president is in New York, and (perhaps) good evidence that the president is not in New York. It is not plausible to suppose that Norman's evidence is such that Norman is justified in closing inquiry, at least not in the right way; the relation between what Norman knows and what Norman's evidence supports is backwards.

Suppose that TSP and radical externalism were both true. Then, Norman (or some other possible subject) knows p and has epistemic justification for not- p . Since he knows p , he is evidentially justified in closing inquiry into whether or not p . What feature of Norman's knowledge explains why he is evidentially justified in closing inquiry? Truth is not the right of thing to explain why Norman is evidentially justified in closing inquiry. The externalist condition, being an external and objective relation in the world, is also not the right kind of thing to explain why

¹¹ Bonjour, *The Structure of Empirical Knowledge*.

Norman is evidentially justified in closing inquiry.¹² That Norman has a relevant belief cannot itself make a difference as whether or not Norman is evidentially justified in closing inquiry. But the only thing left is that Norman has epistemic justification for not-*p*. Norman's epistemic justification might provide him with evidential justification to close inquiry into whether or not *p*, but there is a mismatch between the way that his epistemic justification evidentially justifies him in closing inquiry (assuming that it does), and what he believes. He believes *p*, but his epistemic justification supports not *p*. Norman should not be evidentially justified in closing inquiry, at least not in the manner that the presupposition that he knows that *p* suggests that he should be evidentially justified in doing so. The lesson is that we should reject radical externalism.

In response to his objection, the proponent of TSP and radical externalism might hold that, while justification is not necessary for knowledge, it is necessary that is not the case that one is justified in believing the denial of the true proposition in question. If this were true, then Norman would not know that the president is in New York, even if he believed it, it were true, and the relevant external relations held.

Suppose, though, that what the crystal ball reliably revealed to Norman was that his wife, whom Norman knew to be pregnant, was going to give birth to a boy. Norman is not justified in believing *my wife is not going to give birth to a boy*. Given TSP, he would be evidentially justified in closing inquiry concerning *my wife is going*

¹² In my defense of (7), I will support the claim that the truth condition and the Gettier condition are not the right kind of things to explain why an agent is evidentially justified in closing inquiry.

to give birth to a boy. But he is not evidentially justified in closing inquiry.

Supposing that knowledge is merely true belief and the relevant external relation in the world, there is nothing about the constituents of knowledge that explain why Norman is justified in closing inquiry. Subtract out each constituent that is not a plausible candidate to explain why Norman is evidentially justified in closing inquiry, and you have nothing left.

Premise 5: The Moderate Externalist Worry

These considerations show that the truth of TSP entails the falsity of radical externalism. There is, however, a less radical form of externalism that one might adopt. Suppose that one grants that knowledge requires justification but insists that whether or not one is justified in believing should itself be understood as obtaining when certain external relations obtain. Such externalists face pressure from cases like the new evil demon case.¹³

In the new evil demon case, one imagines that one has a phenomenological twin, who shares one's exact sensory inputs and mental states. The only difference between my twin and I is that my senses are generally veridical, but my twin resides in a world in which an evil demon has thwarted his senses at every turn. When I come to believe that my shoes are in the living room, it is because I remember leaving my shoes are in the living room. When my twin comes to believe that his shoes are in the living room, it is because the demon monkeyed with his experiences such that he had the experience of leaving his shoes in the living room, but did not. Since my twin and I share exactly the same experiences and mental states, we are

¹³ Lehrer and Cohen, "Justification, Truth, and Coherence," 192–193.

justified in believing all the same claims, despite the fact that he knows virtually nothing and I know a lot. After all, assuming that my beliefs are justified, what else should my twin believe except just what I believe? The moderate externalist who insists that justification obtains when the relevant external relation is present faces a problem, because the new evil demon story can be told in ways such that my twin shares all my internal experiences, but lacks whatever external relation to the world that the moderate externalist uses to explain the presence of justification.

I take it that most people, were they to consider this question, would find it intuitive that my twin and I do not differ with respect to our epistemic justification, regardless of the fact that my twin is not (to use one externalist story) forming his beliefs on the basis of reliable belief forming processes. I will consider two ways in which the moderate externalist might reply, and show that both replies entail the falsity of TSP.

One reply the moderate externalist might make is to develop an error theory. This response insists that the intuition that my twin is justified is mistaken, and provides a reason for the widespread error of intuition. As F&M note, such externalists can claim that the intuition that my twin and I do not differ with respect to justification is sensitive to the fact that my twin is free from epistemic blame. People, realizing that my twin is not blameworthy for his beliefs, are hesitant to say that his beliefs are unjustified, which seems to imply culpability. This sensitivity to the fact that my twin is not blameworthy combined with the error of thinking that one is always blameworthy when one believes that which one is not justified in believing causes people to make the mistake of thinking that my twin is justified.

Closing inquiry, however, is an action. We often think people are blameworthy when they stop an investigation at the wrong time. Suppose my twin and I are both considering whether or not we should close inquiry concerning *this children's aspirin doesn't contain any arsenic and is safe to give to my children*. Suppose my child has a bad headache and really wants some pain relief, and this particular aspirin is all that is available. I have some reason to think arsenic has been placed in some bottles of children's aspirin. (My twin has the some experiences, though for him they are not veridical). We both decide to close inquiry. It seems natural to think that, if I am justified in the closure, then so is my twin. Suppose I am justified: Can the moderate externalist explain how we are both justified in closing inquiry? Or must she supply another error theory?

In at least some cases, I don't see how the moderate externalist can make the case that my twin and I are both evidentially justified in closing inquiry. Suppose that I know that the aspirin is safe. Then, it's okay for me to close inquiry (and perhaps to give my children the aspirin). But my twin, according to the moderate externalist, is not justified in believing the aspirin is safe. Shouldn't my twin, if he actually is unjustified in believing that the aspirin is safe, keep checking?

TSP makes the error theory provided by the moderate externalist far less plausible. The original error theory needs it to be the case that my twin is not blameworthy. If TSP is true, then I am justified in closing inquiry. But, assuming that the moderate externalist is correct that my twin lacks justification to believe *this aspirin is safe*, it is hard to see how my twin might be justified in closing inquiry. Closing inquiry is an act, and it is blameworthy to do unjustified acts. So, if both TSP

and the moderate externalist are correct, then my twin should be blameworthy for closing inquiry concerning *this aspirin is safe*. But my twin should be just as justified (or unjustified) as I in the closure of inquiry, since we share identical experiences and mental states.

We see, then, that insisting that my twin is not justified in his beliefs and attempting to develop an error theory to explain why it seems to many people that my twin is justified cannot account for the fact that my twin and I should be similarly justified (or unjustified) in the act of closing inquiry. In order to reconcile TSP and moderate externalism, one needs to not merely develop an error theory, but also explain how it is that my twin and I are equally justified in closing inquiry. A second response to the new evil demon problem, a response developed by Goldman, might help.¹⁴

Goldman distinguishes between weak and strong justification. S's belief is weakly justified if "(1) the cognitive process that produces the belief is unreliable, but (2) S does not believe that the producing process is unreliable, and (3) S neither possess, nor has available to him/her, a reliable way of telling that the process is unreliable."¹⁵ S's belief is strongly justified if and only if "(1) it is produced (or

¹⁴ Goldman, "Strong and Weak Justification."

¹⁵ Ibid., 59. Notice that conditions here are sufficient but not necessary. Further, Goldman adds that one further condition, a condition that involves processes that S believes to be reliable, that he concedes may need to be added to these three conditions. For our purposes, I ignore the possibility of this condition.

sustained) by a sufficiently reliable cognitive process, and (2) that the producing process is reliable is not undermined by S's cognitive state."¹⁶

My phenomenological twin who is being systematically deceived pretty clearly attains weak justification for his beliefs. While there is one mild worry (namely, the worry that my twin *does* have a reliable method available to him that would produce the conclusion that his experiences are not veridical), it seems as though my twin does achieve weak justification for his beliefs. Nonetheless, it is not enough for the moderate externalist to give a plausible case that my twin attains weak justification. By itself, such a response would fail for the same reasons that the previous response (that is, the error theory response) failed. Namely, weak justification, as Goldman concedes, is what a belief attains that allows it to be "faultless, blameless, or non-culpable," while strong justification is needed for a belief to be "well-formed...proper, suitable."¹⁷ Closure of Inquiry, however, is about an action. Suppose my twin stops inquiring into *this aspirin is safe*. Surely it is not enough to say that my twin is justified in closing inquiry merely because his belief that *this aspirin is safe* is not blameworthy. If my twin lacks strong justification for the proposition in question, then his belief is not well-formed; it is improper and

¹⁶ Ibid., 58. Notice that the conditions here are necessary and sufficient. I should also note that Goldman distinguishes between secondary and primary justification by distinguishing between methods that are reliable (secondary justification) and processes that are reliable (primary justification). While this distinction is an important one, I can see no reason why it affects our present considerations. I should also note that the apparent typo in the second condition is found in Goldman's article. I quote it as is, but understand it as not intending to include the "is reliable." I understand the condition as a caveat for internal defeaters.

¹⁷ Ibid., 52–53.

unsuitable. It stretches the bounds of plausibility to claim both that my twin is justified in closing inquiry concerning *this aspirin is safe* and that his belief that *this aspirin is safe* is not a well-formed, proper, or suitable belief.

What the moderate externalist needs is a way of claiming that my twin's beliefs are strongly justified. Recall that Goldman says that S's belief is strongly justified if and only if (1) it is produced (or sustained) by a sufficiently reliable cognitive process, and (2) that the producing process is not undermined by S's cognitive state. In an earlier version of his theory, published in *Epistemology and Cognition* (hereafter, *E&C*), Goldman had the resources to claim that my twin's beliefs were strongly justified.¹⁸ In his earlier theory, Goldman claimed that a belief qualifies as justified "just in case it is produced by processes that conform to some right rules system, " and that a rules system is right just in case it is reliable.¹⁹ In his earlier theory, he clarifies this notion by claiming that the rightness of a rules system is rigid. He says that a rules system "is either right in all possible worlds or wrong in all possible worlds" and that reliability is not determined by the actual world, but by normal worlds, where a normal world is a world that is consistent with our general beliefs about the actual world.²⁰

Under this, earlier, theory (call it the *normal worlds theory*), we might plausibly get the result that my twin is strongly justified (and, hence, it would make sense to claim that he is evidentially justified in closing inquiry in and only those

¹⁸ Goldman, *Epistemology and Cognition*, 107.

¹⁹ Goldman, "Strong and Weak Justification," 61.

²⁰ Ibid.

occasions in which I am evidentially justified in closing inquiry). Unfortunately, as Goldman himself came to realize, the normal worlds theory is implausible. While Goldman lists a number of problems with the normal worlds theory, one particularly difficult problem is sufficient to show its plausibility. Namely, there are cases in which the theory gets the wrong result. There is a possible non-normal world in which people in that world form beliefs on the basis of a process that is very reliable in that world, but is unreliable in normal worlds. Surely the intuitions that attract people to reliabilism should cause us to conclude that people in that world, when they form beliefs in that way, are strongly justified in their beliefs. However, if we don't allow for concept of reliability to be rigidified to normal worlds, it does not appear possible to interpret the notion of strong justification in a way that makes it possible that my twin is strongly justified. The best we can say is that he is weakly justified, and this is not enough to reconcile moderate externalism with TSP. Since we good reason to accept TSP (see chapter two), we should reject moderate externalism. I now turn to a defense of (7).

Premise 7

Recall (7):

7. That p is true and that S 's belief that p is undefeated cannot explain why S is evidentially justified in closing inquiry.

The truth of a proposition is not the kind of thing that can explain why someone is justified in closing inquiry. It seems unhelpful in the extreme to take time considering whether *S 's evidence justifies S in closing inquiry concerning whether or*

not p iff *p* is true. For every *p*, *p* or its denial is true. But there are many propositions for which many of us are not justified in closing inquiry.

Nor could the mere absence or presence of external defeaters for *S*'s belief that *p* explain why *S* is justified in closing inquiry concerning *p*. When I speak of externally undefeated belief, I have in mind the Gettier condition of knowledge. In one of the classic Gettier cases, Smith believes that *the man who will get the job has 10 coins in his pocket*. Smith believes this because he justifiably believes that Jones will get the job, and he knows that Jones has ten coins in his pocket. Smith (not Jones) gets the job, and Smith also has ten coins in his pocket, though Smith had no reason to think that he had ten coins in his pocket.²¹ It turns out that Smith's belief was justified and true, but most people don't think it was knowledge. So, we add the fourth condition. I will follow others in saying that the justified true belief must also be *externally undefeated*.

Defeaters are propositions that defeat support relations. Some defeaters are such that the epistemic agent in question actually considers them; Regarding internal defeaters, it makes sense to say, "If you think you have one, you do."²² These are internal defeaters. In the Gettier case, Smith has a defeater for the proposition that *the man who will get the job has 10 coins in his pocket*, but the defeater is not one that is available to Smith in the relevant sense. His defeater is an external

²¹ Gettier, "Is Justified True Belief Knowledge?"

²² Greco, "Externalism and Defeat." Greco is quoting Jon Kvanvig. The idea here is that internal defeaters are entirely subjective. It should also be noted that they can come as partial defeaters, and they can themselves be defeated by other defeaters. Hence, that *S* has an internal defeater for a proposition is not sufficient for *S* to be unjustified in believing that proposition.

defeater. That a belief be undefeated by external defeaters seems to be one of the fundamental elements of knowledge.²³

Merely the absence of an external defeater is obviously not enough for justified closure of inquiry. Further, if *S is evidentially justified in closing inquiry concerning p iff S's belief that p is externally undefeated* is true, then S would be evidentially justified in closing inquiry on propositions about which she lacks justification to believe. For the reasons given in the last section, that is unacceptable. On the other hand, the presence of epistemic justification, even in cases in which we subsequently learn that an external defeater was in place, is sufficient to justify a closing of inquiry. The jury who finds a defendant guilty on overwhelming evidence should not be criticized merely because it is later learned than they did not know the defendant was guilty because their justification for believing that he was guilty was externally defeated.

Suppose a jury believes *the defendant killed the victim* on the basis of DNA evidence, eyewitness testimony, and a confession. If it was later learned that the defendant did kill the man, but the DNA evidence was planted, the witnesses mistaken, and the confession coerced, it does not follow that the jury was not justified in closing inquiry. In fact, the jury may well have been evidentially justified in closing inquiry.

²³ There is obviously a great deal of controversy regarding how the Gettier problem is best resolved. It is highly controversial to claim that adding freedom from external defeat to the traditional JTB account of knowledge is the best way to handle the problem. Perhaps it isn't. As long as the proper way to handle the Gettier condition involves a fundamental constituent of knowledge that is relevantly external to the agent, the arguments of this section stand; it is not plausible to claim that external facts are the kinds of things that can rob an agent of evidentially justified closure of inquiry.

Nor does a combination of the truth constituent with the externally undefeated constituent explain why a closure of inquiry is evidentially justified. Suppose I believe that *there are is even number of stars in the universe*, and this belief is both true and externally undefeated. I am clearly not evidentially justified in closing inquiry on this proposition – though, supposing I don’t care and it is proper for me not to care, I might well be morally and practically justified in closing inquiry.

The reason that I am not evidentially justified in closing inquiry regarding *there is an even number of stars in the universe* is because I lack epistemic justification in support of the proposition. Add that element, and I am evidentially justified. Further, once that element is added, the presence or absence of the other fundamental constituents does no work. If the jury is epistemically justified in believing *the defendant murdered the victim*, then the jury is evidentially justified in closing inquiry, regardless of the presence or absence of the truth constituent or the Gettier constituent.

Premise 9

Remember what EJ” says:

EJ” S’s evidence justifies S in closing inquiry concerning whether or not *p* iff S’s belief that *p* (or S’s belief that $\sim p$) is epistemically justified.

Remember also what (9) says:

9. If that S is epistemically justified in believing *p* explains why S is evidentially justified in closing inquiry concerning *p*, then EJ” is true.

Notice that (9) is not true merely in virtue of its logical construction. In EJ”, replace *S’s evidence justifies S in closing inquiry concerning whether or not p* with the arbitrary proposition, *c* (for ‘closure’). Replace *S’s belief that p is epistemically*

justified with *j* (for ‘epistemically justified’). *EJ*” reads *c iff j*. (9) seems to claim that if *j* explains why *c* is true, then *c iff j* is true. Can we replace *c* and *j* with propositions in which the antecedent of the conditional is true, but the conclusion is false?

It seems that we can. That *I am my son’s father* (partially) explains why *he is my son*, but it is not true that *he is my son iff I am his father*. Why, then, should we think that *EJ*” follows from the truth (8)?

Let me restate the problem. That *I am Jed’s father* explains *Jed is my son* without *Jed is my son* entailing *I am Jed’s father* shows that the truth of (9) rests not merely on the logical connections. I need to show not only that *S has epistemic justification* explains why *S is evidentially justified to close inquiry*, but also that *S is evidentially justified to close inquiry* entails *S has epistemic justification*. The truth of the latter entailment cannot be demonstrated merely by the logical connections and the nature of explanation.

Luckily, I think it is fairly obvious that *S is evidentially justified to close inquiry concerning p* entails *S has epistemic justification for p (or ~p)*. After all, I stipulated that one is evidentially justified to close inquiry when one is epistemically rational (with the concerns of morality and practical rationality bracketed off) in closing inquiry. The harder question, the one I think that this argument is useful in showing us, is that *S has epistemic justification* entails that *S is evidentially justified to close inquiry*. I take myself to have shown the truth of the hard-to-demonstrate entailment by showing that knowledge entails evidentially justified closure of inquiry, and that (because the other elements can be subtracted without losing the evidential justification to close inquiry) it must do so in virtue of its fundamental

relationship to epistemic justification. If knowledge itself is what does the work of evidentially justifying a closure of inquiry, taking away any of the constituents of knowledge should make a difference as to whether or not one was evidentially justified in closing inquiry.

An Upshot of EJ''

Here is another mark in favor of EJ'': it does much of the theoretical work that a similar principle advanced by F&M does, and it does so without being subject to a certain counterexample to which F&M's principle is vulnerable. In this section, I will articulate the relevant principle that F&M support, the counterexample to which it falls, and show why EJ'' does not have that worry.

Recall that F&M argue for a principle that, while similar to TSP, is significantly stronger. They argue in favor of KJ:

KJ If you know that p , then p is warranted enough to justify you in φ -ing, for any φ .²⁴

After establishing KJ, they think that they can show that the following principle must also be true:

J^KJ If you have knowledge-level justification that p , then p is warranted enough to justify you in φ -ing, for any φ .²⁵

F&M think that not only is J^KJ true, but a bi-conditional version of J^KJ is true²⁶.

²⁴ Greco, "Externalism and Defeat."

²⁵ Ibid.

²⁶ They actually argue for a bi-conditional version of a slightly different principle—a principle they name JJ—but the reasons they have in favor of a bi-conditional JJ also apply to a bi-conditional J^KJ

B-J^KJ You have knowledge-level justification that p iff p is warranted enough to justify you in φ -ing, for any φ .

One of their arguments for J^KJ is the subtraction argument. KJ is the first premise in the subtraction argument. The second premise is that “holding fixed knowledge-level justification while subtracting from knowledge any combination of truth, belief, and being unGettiered makes no difference to whether p is warranted enough to justify.”²⁷ From these two premises alone, J^KJ follows. (The move from J^KJ to B-J^KJ comes later).

In the last chapter, I pointed out that TSP is similar to (though weaker than) KJ. We might also notice that EJ” is a similar to (though weaker than) B-J^KJ.

EJ” S’s evidence justifies S in closing inquiry concerning whether or not p iff S’s belief that p (or S’s belief that $\sim p$) is epistemically justified.

F&M claim that knowledge-level justification is warranted enough to justify one in any action whatsoever. I claim only that epistemic justification (as distinct from alethic justification) evidentially justifies one in a certain type of action, that of closing an inquiry.

While TSP and EJ” entail that knowledge that p and epistemic justification for p are each sufficient for evidentially justified closure of inquiry regarding p , the principles that F&M favor entail that knowledge of p and epistemic justification for p are each warranted enough to justify one in doing any act; they also claim that p warranted enough for one iff no weakness in one’s epistemic position in respect to p stand in the way of p justifying one. Thus, if they are correct, then if one knows that (or has epistemic justification for) p , then if one is not justified in φ -ing, one’s failure

²⁷ Fantl and McGrath, *Knowledge in an Uncertain World*, 99.

to be so justified is not a result of weakness in one's epistemic position with respect to *p*. There is, however, a convincing counterexample to these principles.

Suppose that John knows that *my car is in the driveway*. Then, if F&M are correct, there is no weakness in John's epistemic position with respect to *my car is in the driveway* that should stand in the way of his φ -ing, for any φ . Suppose also someone proposes to place a bet with John: If John is metaphysically certain that *my car is in the driveway*, then John wins. If he is not, then John loses. John is not justified in taking this bet. Does a weakness in his epistemic position regarding *my car is in this driveway* stand in the way of his taking the bet?²⁸ It certainly seems so, but we need to take a closer look at what it means to stand in the way. F&M say this:

We offer no analysis of the intuitive notion of 'standing in the way'. But we do think that, when *Y* does not obtain, the following counterfactual condition is *sufficient* for a subject's position on some dimension *d* to be something that stands in the way of *Y*'s obtaining: whether *Y* obtains can vary with variations in the subject's position on *d*, holding fixed all other factors relevant to whether *Y* obtains. Since we can vary whether or not a 30-year-old is eligible to hold the office of President of the United States by varying her age, holding fixed all other relevant factors, her age stands in the way of her being so eligible.²⁹

In the above story, hold fixed all other factors relevant to whether or not John is justified in taking the bet. Only change the strength of his epistemic position with respect to *p* in the following way: suppose that John possessed infallible knowledge of all true propositions, including *p*. After this change is made, John is justified in taking the bet. Hence, we have an example of a case in which *S* knows *p*, and yet a

²⁸ This counterexample is not my own; I first heard of it in a conversation with Jonathan Kvanvig.

²⁹ Fantl and McGrath, *Knowledge in an Uncertain World*, 67.

weakness in S 's strength of epistemic position with respect to p stands in the way of S 's φ -ing.

Since TSP and EJ'' are restricted to closures of inquiry, they do not fall prey to this counterexample. Further, TSP and EJ'' only speak of when one is evidentially justified to close inquiry. This qualification makes TSP and EJ'' less informative until I can say more about the relationship between evidential justification and full (or all-things-considered) justification. In the next chapter, that is my task.

CHAPTER FIVE

Fully Justified Closure of Inquiry

In chapter one, I argued that any successful theory of justified closure of inquiry must include a role for the epistemic domain. In chapters two and three, I argued that EJ'' is true; one is evidentially justified in closing inquiry if and only if one's relevant belief is epistemically justified. While these conclusions are helpful, the central question of this investigation is this: What justifies closure of inquiry? The truth of EJ'' does not itself fully answer the question, because EJ'' only specifies the necessary and sufficient conditions under which one is evidentially justified in closing inquiry. Suppose that Sam is thinking about closing his inquiry regarding some conclusion. He wants to know if he is justified in doing so. Suppose further that he is evidentially justified, but not practically justified (or, alternatively, not morally justified) in closing inquiry. I tell him, "Sam, you are evidentially but not practically (or morally) justified in closing inquiry." I think he might reply, "Thanks for the information, but I don't care about that. I just want to know if I am justified in closing inquiry."

I call the kind of justification Sam is asking about *full justification*. I take it that this notion is fairly intuitive, even though I've only stated it loosely. Full justification is all-things-considered justification; it is the justification one has after all normative constraints have been successfully considered in full, and the judgment is that one is permitted to act. My central question is this: When does full

justification to close inquiry obtain? Since it is possible for one to be morally prohibited from closing inquiry in situations in which one has evidential justification to close inquiry, evidential justification does not entail full justification. Hence, I need to say more about full justification and when it obtains. EJ'' is not enough.

In articulating when full justification obtains, the most difficult issue is the relationship between evidential and practical justification to close inquiry. There are two difficult questions regarding this relationship. First, can one be evidentially justified in closing inquiry, but not practically justified? I do not attempt to argue one way or the other in regards to this question. Instead, I look at the effect that a negative answer has on a theory of justified closure of inquiry and the effect that a positive answer has on such a theory. Supposing that it is possible for one to be evidentially but not practically justified in closing an inquiry, a second central question arises; if one finds oneself in a situation in which a closure is evidentially but not practically justified, is one ever fully justified in closing inquiry? My answer to this second question is dependent upon the working theory of practical rationality. If practical rationality is about the maximization of expected utility, then it is possible for one to be fully justified in closing inquiry even when it is not practically rational for one to do so; if practical rationality is itself dependent upon epistemic rationality (in the manner I discussed in chapter one), then one who is not morally obligated to close inquiry is never fully justified in closing inquiry unless it is practically rational for her to do so.

My aim is to show that the correct theory of justified closure of inquiry is non-reductive or, if it can be reduced, it is reduced by the elimination of the practical domain. As to which of those two disjuncts makes the claim true, I take no stand, although I favor a non-reductive account. If one cannot be evidentially justified without being practically justified, then the most plausible account of justified closure of inquiry can be stated without mention of practical justification. If one can be evidentially justified without being practically justified, then the correct account of justified closure of inquiry is non-reductive. In defense of this thesis, I need first to provide a better explication of the non-reductive theory that I have in mind. After explicating the theory, I argue in favor of my disjunctive thesis.

A Non-Reductive Theory of Justified Closure of Inquiry

A theory of justified closure of inquiry is non-reductive if it explains the justificatory status of the closure in terms that make fundamental use of each of the three major domains of normativity: morality, practical rationality, and epistemic normativity. Obviously, there are very many possible theories that are non-reductive in this sense; most of them are not very plausible or interesting. In trying to get a better sense of what an interesting and plausible non-reductive theory looks like, it is helpful to consider the possible ways that the three major domains of normativity might relate to one another:

1.	M	PR	EJ
2.	M	PR	EJ
3.	M	PR	EJ
4.	M	PR	EJ
5.	<i>M</i>	PR	EJ
6.	<i>M</i>	<i>PR</i>	EJ
7.	<i>M</i>	<i>PR</i>	<i>EJ</i>
8.	<i>M</i>	PR	<i>EJ</i>

The ‘M’ in the above list represents morality, ‘PR’ represents practical rationality, and ‘EJ’ represents epistemic justification. Black letters represent occasions when the given domain justifies closing some given inquiry, and the red (or italicized) letters represent occasions when the given domain does not justify closing inquiry. A plausible non-reductive theory needs to have something to say about whether or not full justification to close inquiry obtains in each of these eight scenarios.

Scenario 1 is simple. If every domain justifies closure of inquiry for S, then S is justified in closing inquiry. Scenario 7 is simple for the same reasons, but in the opposite direction.

Scenarios 5 through 8 do not provide S with full justification to close inquiry. One is never fully justified in doing anything that one morally ought not do. Not only is it incredible counter-intuitive to claim that one can be fully justified in doing that which one morally ought not do, but moral oughtness is the paradigm normative concept. As the paradigm normative concept, it is both unfitting and exceedingly strange to claim that one is fully justified in doing that which one morally ought not do. The more difficult questions arise when considering scenario 2, 3, or 4.

In considering the possibility of these three scenarios, and whether full justification obtains in the event that any of them occur, there is a distinction I need to make between different species of the relevant scenarios:

1A.	M	PR	EJ
1B.	<u>M</u>	PR	EJ
2A.	M	PR	EJ
2B.	<u>M</u>	PR	EJ
3A.	M	PR	EJ
3B.	<u>M</u>	PR	EJ
4A.	M	PR	EJ
4B.	<u>M</u>	PR	EJ

Scenarios 1 through 4 can each be further divided into those cases in which one is both morally permitted to close inquiry and morally obligated to close inquiry (1B, 2B, 3B, and 4B) and those cases in which one is morally permitted to close inquiry, but not morally obligated (1A, 2A, 3A, and 4A). The blue (or underlined) letters in the above list represent moral obligation to close inquiry, while the black letters represent justification, but no obligation. Since one who is morally obligated to close inquiry is morally justified in closing inquiry, the cases of obligation are a species of the cases of moral justification.

Consider scenario 3B. We might imagine that I have been inquiring into whether or not, say, my shoes are on the porch. As I look on the porch and see and notice that my shoes are on the porch, I learn that a moral horror will occur unless I keep looking for my shoes. I can't just pretend to look for my shoes, mind you. In order to avoid the moral horror, I must keep inquiring. I could, for example, begin looking at the shoes more closely, looking for evidence that these shoes merely look very much like my shoes. It would seem that in at least some possible cases in which I must keep inquiring in order to avoid a moral horror, I am morally obligated to keep inquiring. It is important to note that this might be the case for me even if I don't at all care about protecting the world from this moral horror. In such a case, it might be practically rational for me to close inquiry, and I would be evidentially

justified in closing inquiry; but I would still not be fully justified in closing inquiry, because one is never fully justified in doing that which is morally impermissible for one.

Situations that fall under scenario 3A are more difficult to diagnose. The tension here is that one is morally permitted to close inquiry (perhaps causing it to seem odd to diagnose the case as one in which one does not have full justification to close inquiry), but doing so is not practically rational (perhaps causing it to seem odd to diagnosis the case as one in which one does have full justification to close inquiry). When discussing other kinds of acts than closures of inquiry, it often does seem correct to say that one ought not do something (where what we have in mind is that they really ought not do it, all things considered) even if it would not be morally impermissible for them to do so. Suppose I am playing basketball, and I very much care about winning. I ought not let my opponent score every possession. Given my practical concerns, all things considered, I ought to try and play good defense. I don't see why acts of closure should be any different. In scenario 3A, as in scenario 3B, one is not fully justified in closing inquiry.

If scenarios 2 and 4 are possible, then evidential and practical justification to close inquiry can come apart. I have said that I don't know whether or not one can be evidentially justified in closing inquiry, but not practically justified; what I mean by this is that I take no stand on whether or not scenario 2 is possible.

Scenario 4, however, is clearly possible. When I was a kid, I (for a short time) cared about whether or not Corvettes or Porsches were faster. I entered inquiry, but soon closed inquiry after coming to not care about the answer. I did not have

evidential justification to close inquiry regarding any of the relevant propositions, but I was (nonetheless) fully justified in closing inquiry; after all, I didn't care about the answer anymore, and I wasn't morally required to investigate. It seems that when one finds oneself in scenario 4, one is justified in closing inquiry, but one does not have epistemic justification in favor of the relevant proposition.

Notice, however, that it is possible for one to be alethically justified in believing the relevant proposition. Given EJ'', a lack of evidential justification entails a lack of epistemic justification, but it does not entail a lack of alethic justification.

To summarize, one has full justification to close inquiry whenever one is morally obligated to close inquiry. In those cases in which one is morally permitted, but not morally obligated, to close inquiry, then one is fully justified in closing inquiry if the following is also true: 1) one is practically and evidentially justified to close inquiry, or 2) one is practically justified but not evidentially justified. In the latter case, then if one is justified in believing *p* (where *p* is the proposition into which one is, or was, inquiring) one merely has alethic justification to believe *p*. If one is morally permitted to close inquiry but lacks practical and evidential justification to do so, then one is not fully justified to close inquiry.

I will now discuss scenario 2A. If scenario 2A is possible, then the correct account of justified closure of inquiry is non-reductive. If scenario 2A is not possible, then the correct theory of justified closure of inquiry can be stated without mention of practical rationality.

Evidential Justification without Practical Justification

Some people seem to hold that the fact that S has epistemic justification for *p* entails that S is practically rational in closing inquiry regarding *p*.¹ After all, at least some of the force behind the intuition that we do not know that *my lottery ticket is a loser* is because we recognize that we are not practically justified in acting on that knowledge. “If I knew it,” one might think, “then I would be *justified* in not checking anymore. But I’m not justified in not checking, so I must not know it.” I draw attention to the word justified in the previous sentence, because it is crucial that we know what kind of justification we have in mind when we consider such an argument. In this chapter, I am concerned with full justification. If the argument I just presented is intended to deal with full justification, then the phrase, “If I knew it, I would be justified in not checking anymore” would be false, because there are cases in which one can know *p* without being justified in failing to check into *p*; namely, in this cases in which one is morally obligated to keep checking into *p*.

The lottery case, though, is not the kind of case in which one is morally obligated to keep checking. It seems fair, then, to understand the argument like this, “Since I’m not morally obligated to check to see if my ticket is a loser, than if I knew that my ticket was a loser, then I would be fully justified in not checking. But I’m not fully justified in failing to check. So, I must not know.” To the extent that the premises of this argument have intuitive appeal (and they do, for me at least), then there is reason to think that evidential justification to close inquiry entails practical justification to close inquiry. If it is true that evidential justification entails practical

¹ Fantl and McGrath, *Knowledge in an Uncertain World*, 96–100.

justification, then situations like scenario 2A are impossible; a theory of justified closure of inquiry could be true while ignoring such scenarios. The resulting theory would be partially reductive, because it could be articulated without any mention of practical normativity.

However, there are good reasons to think that evidential justification does not entail practical justification. For one, scenario 2B (the situation in which one of morally obligated to close inquiry, even though doing so is practically unjustified and evidentially justified) seems possible. If evidential justification fails to entail practical justification given the presence of moral obligation, why think evidential justification entails practical justification in the absence of moral obligation?

Secondly, the claim that evidential justification entails practical justification makes it more difficult (though perhaps not impossible) to defend the theory of justified closure of inquiry from cases like Brown's surgeon case. If one admits the possibility that one can be evidentially justified without being practically justified, then one can allow the result that the surgeon knows which kidney to remove and that the surgeon is not fully justified in refraining from checking the records before removing the kidney. It is more difficult to admit this possibility while holding that evidential justification entails practical justification.²

Suppose, then, that cases like Brown's surgeon case convince one that it is possible for one to have evidential justification without having practical

² In order to admit the result that the surgeon is not fully justified in failing to check the records even while holding that evidential justification entails practical justification, one would need to hold that the surgeon is morally obligated to check the records, or that some other normative domain is at work; In the next section, I'll consider the concern that more normative domains than the three that I consider play into justified closure of inquiry.

justification. In that case, it seems that my thesis is in trouble. If the only scenarios in which one is justified in closing inquiry are those scenarios in which one is practically rational in doing so or morally obligated to do so, then it seems we can give an account of justified closure of inquiry without appealing to epistemic normativity. Recall an account that I rejected in chapter one:

(PR \wedge MJ) \vee MO: S is justified in closing inquiry concerning whether or not p iff
 a) it is practically rational for S to close inquiry into whether or
 not p and S is morally justified in closing inquiry into whether
 or not p , or b) S is morally obligated to close inquiry.

In chapter one, I pointed out that such an account is only reductive if practical rationality is not itself fundamentally dependent on the epistemic domain, and I argued that there is a plausible account of practical rationality that is fundamentally dependent on epistemic justification. For example, it seems plausible that an act is practically rational (relative to some end) for one if and only if one is justified of believing of the act that it is as good as any other acts in regards to achieving the relevant end.³ Under this conception of practical rationality, the claim that one is not fully justified in closing inquiry if one is not practically justified in closing inquiry is equivalent to the claim that one is not fully justified in closing inquiry if one is not justified in believing of the closure that it is as good as any other means in regards to achieving the relevant end. You may also recall that whether a means is as good as any other means is sensitive to other ends one has, perhaps to

³ The careful reader may notice that, given the distinction between alethic and epistemic justification, one can be justified in believing of the closure that it is as good as an other available means without having epistemic justification for the belief. Practical justification, then, would not be dependent on epistemic justification, but on alethic justification. But that's okay, because alethic justification resides in the epistemic domain.

the totality of one's ends. If this is the working account of practical rationality, then the theory under consideration retains its status as non-reductive, because it makes fundamental use of all three of the relevant domains of normativity.

While the above account is plausible, it parts ways with the most popular way of understanding practical rationality. Under this conception of practical rationality, "rational behaviour is maximizing behaviour relative to one's preferences (and beliefs)."⁴ According to the utility maximization model of practical rationality, a given act is the practically rational thing for S to do if and only if performing the act maximizes S's expected utility. Expected utility is a function of what S believes about the probability of the outcomes of the actions and S's preferences in regards to those potential outcomes. One maximizes behavior relative to one's preferences and beliefs (ie, is practically rational) when one does the act with the greatest expected utility. Those who affirm an expected utility maximization model of practical rationality might hold that my thesis is false, because a correct account of justified closure of inquiry can be articulated using only the notion of moral justification and practical rationality. While I argued against just this conclusion in chapter one, I'd like to give it a more sustained treatment here – a treatment that enjoys the benefit of a clearer existing distinction of the relationship between epistemic justification and justified closure of inquiry (as described by EJ").

I argue, first, that the maximization of expected utility model of practical rationality is false. For those who are not convinced by my arguments that this model is false, I argue that, if it is true, then in at least some situations like scenario

⁴ Sowden, "The Inadequacy of Bayesian Decision Theory," 294.

2A, one is fully justified in closing inquiry. In order to distinguish between those cases of scenario 2A in which one is fully justified in closing inquiry and those cases in which one is not, one needs to appeal to epistemic normativity. In other words, either the maximization theory is false or it is true. If it is false, then a conception of practical rationality that is dependent on the epistemically normative (the truth of which would lead the correct theory of justified closure of inquiry to be non-reductive) is incrementally advanced. If it is true, then one must explain those cases of scenario 2A in which closure of inquiry is fully justified by making appeal to what one has epistemic reason to believe. In either case, a non-reductive account is advanced.

Practically Rationality and Utility Maximization

There is a simple argument against the maximization of expected utility model of practical rationality. Many of the acts that we do everyday do not maximize our subjective expected utility and yet are practically rational. I just drove from my house to work, which is on the other side of town. I don't know how many routes are such that they are reasonable candidates for being the route that, were I were to perform the expected utility calculations, would win the award for maximizing my expected utility, but it is at least in the dozens; I don't know if the route I took would win that award or not. I can give rough estimates regarding the utility of some of the different routes. For example, I might have driven down Park Lake and along the river. I enjoy driving in pretty places and in lighter traffic. That route would have taken longer, however, and getting to work quickly has some positive utility. On some occasions when I am late for an important meeting, it may

be obvious that the utility of getting to work more quickly obviously outweighs the utility of enjoying a pretty drive. On other occasions, however, the percentages and utilities involved may be psychologically obscure. There might still be a fact of the matter, even though it is not obvious to me what the fact of the matter is, despite it being the case that the fact of the matter pertains to my own psychological state. Furthermore, even if it is obvious, for example, that the speedier route maximizes my expected utility, there are a number of small deviations I can make along that route, and it is not obvious to me what the outcome of an expected utility calculation among those smaller deviations would be. It is not even obvious that the facts about my degrees of belief and utility assignments are sufficiently fine-grained to perform such a calculation.

While I can't be sure that the route I ultimately took maximized my expected utility, I do know that I was practically rational in my selection. When I set my alarm last night, I chose 6:15. I could have chosen 6:14, 6:13, or a number of other reasonable (for me) times. I don't know which of those times maximizes my expected utility, but I do know that I was practically rational in choosing 6:15. When I trimmed the bushes yesterday, I chose how long to trim the branches. I might have chosen a little shorter, or a little longer. I don't know if the length I chose maximized expected utility for me or not, but I do know that I was practically rational in trimming them the length that I trimmed them.

We might put the objection like this: If that an act is practically rationality for an agent requires that the act maximize her expected utility, then (for example) I was irrational in setting my alarm for 6:15. But I was not irrational in setting my

alarm for 6:15. Therefore, the antecedent is false. Call this the demandingness objection.⁵

It might be objected that, as described, these are not cases in which practical rationality comes apart from expected utility. Rather, these are cases in which I know that the act was rational, but I don't know that the act maximized my expected utility. That seems true. But it seems highly plausible to me that, in at least some such cases, the act that I describe actually fails to maximize expected utility. Suppose that, if I were to carefully and painstakingly analyze my beliefs and preferences, I would see that setting the alarm for 6:14 has an ever-so-slightly higher expected utility than setting the alarm for 6:15. So what? The difference is very small and the calculation difficult (even supposing that it is possible) to perform. Setting the alarm for 6:15 was practically rational in such a case, regardless of the fact that it does not maximize my expected utility.

There are two ways proponents of the maximization of expected utility model might respond to the demandingness objection. First, they might insist that the theory is purely descriptive, not normative. When a proponent of a descriptive version of the maximization models says that one is behaving irrationally, what they say should not be understood as entailing that the relevant agent is behaving in a manner that is blameworthy, bad, or in any sense on the wrong side of any normative tracks. There seems to be something to this response to the demandingness objection. After all, there is something noteworthy about the fact

⁵ Sobel, "The Impotence of the Demandingness Objection," 2. Notice the similarity between the objection that I am lodging here and the demandingness objection to consequentialism.

that an act maximizes one's expected utility. However, such a theory has nothing to say in regards to when one is fully justified, because the concept that I have in mind is normative. We can and often do blame people (or, even when we deem culpability is not present, negatively judge their action in some normative respect) for closing inquiry when they are not fully justified in doing so. Hence, a purely descriptive theory of a maximization model of practical rationality, while interesting, is not the right kind of theory to be employed in our present endeavor.

A second response to the demandingness objection is to move to a satisficing model. Perhaps it is not necessary that one maximize expected utility in order to be practically rational, but only that one achieves some acceptable level of expected utility. If an act achieves at least some minimum threshold of expected utility for one, then one is practically rational in acting. Perhaps, then, while setting my alarm for 6:13 would have maximized my expected utility (and, hence, under a straightforward maximization model, I was practically irrational in setting the alarm for 6:15), setting the alarm for 6:15 achieved enough expected utility to pass the threshold; it satisfied.

Under one satisficing model, one's expected utility is maximized by not attempting to maximize the utility of every act. I could either inquire into whether or not setting my alarm at 6:15 maximizes my expected utility, or not. It happens that refusing to inquire maximizes my expected utility. Hence, under this satisficing model, I am better off by choosing to set my alarm at 6:15, even if doing so fails to maximize expected utility, because the cost of doing the calculation is high enough

to offset the benefits of getting it right. Call this Strategic Practical Satisficing.⁶ Other satisficer's think that even when an agent knows that a certain act fails to maximize utility, she is still rational in doing the act as long as it exceeds the relevant threshold of utility expectation.⁷ Call this Blatant Practical Satisficing.

Strategic Practical Satisficing is insufficient as a response to the demandingness problem. Suppose I am deciding whether or not I should set the alarm at 6:15. Suppose it happens to be the case that performing the inquiry to see whether or not setting the alarm at 6:15 is practically rational (i.e., that is maximizes my expected utility) has a lower expected utility for me than not performing the inquiry. Then, according to a straightforward maximization model, it is not practically rational for me to perform the inquiry, even if failing to do so has the result of my failing to maximize utility with respect to the time that I set the alarm. Hence, according to the straightforward maximization model, I was irrational in setting the alarm for 6:15 (since doing so did not maximize my expected utility), but was rational in not inquiring into whether or not the act was rational. This is the wrong result, since it seems to follow that a great number of acts that are rational get judged as irrational by the maximization model. In this case, Strategic Practical Satisficing gets the right result; both actions (the act of setting the alarm for 6:15, and the act of closing inquiring into whether or not setting the alarm for 6:15 maximized my expected utility) were practically rational. However, the verdict of

⁶ Mulgan, "Slote's Satisficing Consequentialism," 121–122.

⁷ Slote, *Beyond Optimizing: A Study of Rational Choice*.

Strategic Practical Satisficing gets it wrong in other cases, in a very similar way to how the straightforward maximization model gets in wrong in this case.

Modify the above case such that performing the inquiry to see whether or not setting the alarm at 6:15 maximizes my utility has a higher expected utility for me than not performing the inquiry. However, suppose that the preferences and beliefs involved with inquiring into whether or not I should inquire about setting the alarm at 6:15 are just as subtle and complex as the inquiry into whether setting the alarm at 6:15 maximizes my expected utility; how should I know whether or not I should perform the second level inquiry? Strategic Practical Satisficing (unless it is understood purely descriptively) judges that I err in all such cases (ie, in all cases in which performing the second level inquiry maximizes my expected utility) by not performing the inquiry; but surely that is false. Just as there times when I am practically rational in setting my alarm for a certain time, even if a slightly different time would maximize my expected utility, surely there are times when I am practically rational in not inquiring into whether or not *setting my alarm to time t maximizes my expected utility*, even in some cases in which performing the inquiry has a higher expected utility than not performing the inquiry. The reason for abandoning straightforward maximization—that it is subject to the demandingness objection—is undercut; call this the meta-level demandingness objection.

Blatant Practical Satisficing, if it is true, might initially seem more like a rejection of maximization theory than a version of it. Since Blatant Practical Satisficing says that one can sometimes be practically rational in closing inquiry even when one knows that continuing to inquire has a higher expected utility, it

might be hard to see how the affirmation of the theory is sometimes viewed as an affirmation of a general maximization schema. A second distinction in types of satisficing theories makes it easier to see why Blatant Practical Satisficing can be a maximization theory. This second distinction, articulated by Hurka, is between absolute level satisficing and comparative satisficing. An absolute level view says that an action must achieve some absolute level of expected utility in order to be rational, but need not maximize beyond that absolute threshold. Of course, if an absolute level satisficing theory is true, then if an agent has no practically rational option available, then the agent could not possibly act in a practically rational manner. Under a comparative satisficing theory, if no available action meets the absolute threshold deemed necessary for the rationality of an action, then an act is rational so long as it achieves a higher expected utility than all the other available actions.⁸ Blatant Practical Satisficing is a maximization theory, then, because it says that rationality demands that one either maximize one's utility or reach a certain threshold. If an act reaches that threshold, then (perhaps) one can choose that act even if one knows that a different action would achieve an even higher utility.

The problem with Blatant Practical Satisficing as a response to the demandingness objection is that it is unmotivated. The appeal of a theory of practical rationality that affirms that rational agents are always expected utility maximizers is the intuitiveness of the idea that one can never rationally and knowingly choose the worse of two options. If one rejects that idea, as the

⁸ Hurka, "Two Kinds of Satisficing," 107–108.

proponent of Blatant Practical Satisficing must, then there is no adequate reason to promote a maximization theory.

For these reasons, both of the best attempts to defend the maximization model from the demandingness objection fail. Either the maximization theory is merely descriptive (in which case it is unhelpful in the discussion of a normative theory of justified closure of inquiry), or it returns the wrong verdict in many instances in which someone has clearly acted in a manner that is practically rational. Strategic Practical Satisficing is subject to a meta-level demandingness objection, and Blatant Practical Satisficing should be unmotivated for anyone who wants to defend a maximization model. Given the failure of a maximization model of practical rationality, a theory of practical rationality that is itself dependent on the epistemic domain is (incrementally) advanced. Given the truth of such a theory, a correct theory of justified closure of inquiry is non-reductive even if scenario 2A is impossible.

Despite these objections to a utility maximization model of practical rationality, many people think that practical rationality is about utility maximization. I'd rather a defense of my thesis not rely on the falsity of all utility maximization models of practical rationality. So let's suppose that a given act is the practically rational thing for S to do if and only if performing the act maximizes S's expected utility, or that some other satisficing theory of practical rationality is true. If this is the case, situations that fit scenario 2A are possible, and in some such situations one is fully justified in closing inquiry.

In chapter one, I used Lara Buchak's idea of risk avoidant agents to argue that the practically rational agent need not always maximize expected utility. Here, I will simply grant that an act is not practically rational for one unless it maximizes expected utility (or, at least, reaches some absolute utility threshold) for one.

Suppose, for example, that S is considering whether or not to close inquiry. Suppose that doing so will not maximize S's expected utility or meet the required absolute expected utility threshold (i.e., given our current assumptions, is not practically rational), but that S is justified in believing that closing inquiry is as good as any other means toward her ends (we might even suppose that S is unjustified in believing that closing inquiry is not the best means given the totality of her ends, or that S is epistemically obligated to believe that closing inquiry is the best means to her end). Suppose also that closing inquiry is morally justified and evidentially justified. In such a case, S is certainly fully justified in closing inquiry, regardless of the fact that closing inquiry is not practically rational for S. If it is true that an act is the practically rational thing for S to do if and only if performing the act maximizes S's expected utility, then situations like 2A are possible, and (in some of them), one has full justification to close inquiry. Hence, an acceptable theory of justified closure of inquiry would need to address those occasions of full justification, and doing so would require the theory to make fundamental use of the epistemic domain.

For example, it might be the case that one is worried that any defeaters one will find in continued inquiry will be misleading defeaters. Suppose that one is morally and evidentially justified in closing inquiry, but closing inquiry does not maximize one's expected utility because a cost-free inquiry is available.

Nonetheless, one is justified in believing that closing inquiry is the best means to one's ends, perhaps because one is justified in believing that further inquiry will reveal only misleading defeaters. One would be fully justified in closing inquiry in such a case, and an adequate theory of justified closure of inquiry would need to include it. Since this addition makes fundamental use of epistemic normativity, the theory in which it was included would no longer be reductive.

Conventions, Etiquette, and Aesthetic Normativity

Here is a concern about the non-reductive theory that I have advanced: it's too simplistic. Practical rationality, morality, and epistemic normativity are not the only domains of normativity. There are many more, such as societal conventions, etiquette, and aesthetic standards. If any of these or other domains of normativity affect full justification, then it might not be true to say (to pick one of the situations in which full justification obtains) that one has full justification to close inquiry anytime one is epistemically justified in closing inquiry, is not morally prohibited from closing inquiry, and is practically rational in closing inquiry.

Suppose, for example, that closing inquiry would breach some rule of etiquette or convention. Such cases are hard to imagine, but they seem possible. For example, suppose that etiquette demands that one not eat one's food until one first turns to see if the host of the party has already taken a bite. Suppose that one knows that the host has taken a bite, and that one's ends are not better served by checking to see if the host has taken a bite. Nonetheless, one may not be justified in failing to check, since etiquette demands it. Given the complexity of the

relationships between different normative domains, the non-reductive theory that I favor should be seen merely as a helpful and important heuristic.

The use of cognitive tools or shortcuts is commonplace. For example, it is well documented that humans are not natural probabilistic thinkers. For this reason, they often use shortcuts to help with their probabilistic decisions making. These tools, while often helpful, can sometimes provide the wrong answer for those who make use of them. For example, people often ignore the importance of sample size when asked about the probability of two classes (of differing sizes) straying from their most likely distributions. This mistake is often blamed on the common use of the representation heuristic, a heuristic in which people assess the likelihood of a sample result from what is representative from some other salient class.⁹ The use of such heuristics, while often helpful in daily decision-making, can lead to mistakes. It might be that my statement of when full justification obtains is actually a reliable heuristic, albeit one that can err in certain unusual circumstances.

Conclusion

I began my dissertation with this central question: When is closure of inquiry justified? I began by showing that different attempts to reduce the normative domains have failed, and that an epistemic component is essential to any attempt to articulate when justified closure of inquiry obtains. In chapter two and three, I showed that evidential justification for one to close inquiry into a proposition rises or falls with one's epistemic justification in regards to the proposition. In chapter four, I finally answered the question. When is one (fully)

⁹ Tversky and Kahneman, "Judgment Under Uncertainty," 1124.

justified in closing inquiry? On a maximization model of practical rationality, one is fully justified to close inquiry whenever one is morally obligated to close inquiry, or one is practically and morally justified in closing inquiry, or one is morally and evidentially justified in closing inquiry and one is justified in believing that closing inquiry is as good as any other means to the achievement of one's ends (perhaps because one is justifiably worried that any defeaters one finds will be misleading defeaters). On a model of practical rationality in which an act is practically rational for one only if one is justified in believing that the act is as good as any other means toward the achievement of one's ends, then one is fully justified in closing inquiry whenever one is morally obligated to close inquiry, or one is morally and practically justified to close inquiry.

BIBLIOGRAPHY

- Adler, Jonathan. "Reconciling Open-Mindedness and Belief." *Theory and Research in Education* 2, no. 2 (July 1, 2004): 127–142.
- Allais, Maurice. "Criticisms of the Postulates and Axioms of the American School." In *Rationality in Action: Contemporary Approaches*, edited by Paul K. Moser. Cambridge University Press, 1990.
- Anscombe, G. E. M. *Intention*. 2nd ed. Harvard University Press, 2000.
- Ayer, Alfred J. *Language, Truth and Logic*. 2nd ed. Dover Publications, 1952.
- Baier, Kurt. "Rationality and Morality." *Erkenntnis: An International Journal of Analytic Philosophy* 11 (1977): 197–223.
- Bishop, John. "Faith as Doxastic Venture." *Religious Studies* 38, no. 4 (December 1, 2002): 471–487.
- Blackburn, Simon. *Essays in Quasi-realism*. Oxford University Press, 1993.
- BonJour, Laurence. *The Structure of Empirical Knowledge*. Harvard University Press, 1985.
- Bradley, Ben. "Against Satisficing Consequentialism." *Utilitas: A Journal of Utilitarian Studies* 18, no. 2 (June 1, 2006): 97–108.
- Brown, Jessica. "Subject-Sensitive Invariantism and the Knowledge Norm for Practical Reasoning." *Noûs* 42, no. 2 (June 1, 2008): 167–189.
- Buchak, Lara. "Instrumental Rationality, Epistemic Rationality, and Evidence-Gathering." *Nous-Supplement: Philosophical Perspectives* 24 (January 1, 2010): 85–120.
- Chisholm, Roderick M. "Firth and the Ethics of Belief." *Philosophy and Phenomenological Research* 51, no. 1 (March 1, 1991): 119–128.
- Christensen, David. "Diachronic Coherence Versus Epistemic Impartiality." *The Philosophical Review* 109, no. 3 (July 1, 2000): 349–371.
- Conee, Earl, and Richard Feldman. *Evidentialism: Essays in Epistemology*. Oxford University Press, USA, 2004.

- Cooke, Elizabeth. "Transcendental Hope: Peirce, Hookway, and Pihlström on the Conditions for Inquiry." *Transactions of the Charles S. Peirce Society* 41, no. 3 (July 1, 2005): 651–674.
- DeRose, Keith. *The Case for Contextualism: Knowledge, Skepticism, and Context, Vol. 1*. Oxford University Press, USA, 2009.
- Dewey, John. *Democracy and Education: Complete and Unabridged*. Unabridged. CreateSpace, 2009.
- Dougherty, Trent. *Evidentialism and Its Discontents*. Oxford University Press, 2011.
- Dougherty, Trent. "Reducing Responsibility: An Evidentialist Account of Epistemic Blame." *European Journal of Philosophy* (n.d.).
<http://onlinelibrary.wiley.com/doi/10.1111/j.1468-0378.2010.00422.x/abstract>.
- Dreier, James. "Rational Preference: Decision Theory as a Theory of Practical Rationality." *Theory and Decision* 40 (May 1996): 249–276.
- Fantl, Jeremy, and Matthew McGrath. "Evidence, Pragmatics, and Justification." *Philosophical Review* 111, no. 1 (January 1, 2002): 67–94.
- Fantl, Jeremy, and Matthew McGrath. *Knowledge in an Uncertain World*. Oxford University Press, USA, 2009.
- Fantl, Jeremy, and Matthew McGrath. *The Routledge companion to epistemology*. Edited by Sven Bernecker and Duncan Pritchard. New York: Routledge, 2011.
- Feldman, Richard. *Epistemology*. 1st ed. Prentice Hall, 2002.
- Foley, Richard. *The Theory of Epistemic Rationality*. Cambridge, Mass: Harvard University Press, 1987.
- Foley, Richard. "The Foundational Role of Epistemology in a General Theory of Rationality." In *Virtue Epistemology: Essays in Epistemic Virtue and Responsibility*, edited by Abrol Fairweather and Linda Zagzebski. Oxford University Press, USA, 2001.
- Fricker, Miranda. *Epistemic Injustice*. Oxford University Press, 2007.
<http://www.oxfordscholarship.com.ezproxy.baylor.edu/view/10.1093/acprof:oso/9780198237907.001.0001/acprof-9780198237907>.
- Geach, Peter Thomas. *Logic Matters*. University of California Press, 1972.

- Gibbard, Allan. *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Oxford University Press, 1992.
- Gibbard, Allan. *Thinking How to Live*. Harvard University Press, 2003.
- Goldman, Alvin I. *Epistemology and Cognition*. Harvard University Press, 1986.
- Goldman, Alvin I. "Strong and Weak Justification." *Philosophical Perspectives* 2 (January 1, 1988): 51–69.
- Good, I J. "On The Principle Of Total Evidence." *British Journal for the Philosophy of Science* 17 (February 1, 1967): 319–321.
- Greco, John. "Further Thoughts on Agent Reliabilism: Replies to Cohen, Geivett, Kvanvig, and Schmitt and Lahroodi." *Philosophy and Phenomenological Research* 66, no. 2 (March 1, 2003): 466–480.
- Haddock, Adrian, Alan Millar, Alan Millar (Ph. D.), and Duncan Pritchard. *Epistemic Value*. Oxford University Press, 2009.
- Hare, William. *In Defence of Open-mindedness*. McGill-Queen's Press - MQUP, 1985.
- Hawthorne, John. *Knowledge and Lotteries*. Oxford : Oxford ; New York: Clarendon ; Oxford University Press, 2004.
- Hetherington, Stephen Cade. *Epistemology Futures*. Oxford University Press, 2006.
- Hickman, Larry A., and Thomas M. Alexander. *The Essential Dewey, Volume 1: Pragmatism, Education, Democracy*. Indiana University Press, 1998.
- Hintikka, Jaakko, Vincent F. Hendricks, and John Symons. *Knowledge and Belief - An Introduction to the Logic of the Two Notions*. College Publications, 2005.
- Hookway, Christopher. *Truth, Rationality, and Pragmatism: Themes from Peirce*. Oxford University Press, 2002.
- Hubin, Donald C. "The Groundless Normativity of Instrumental Rationality." *The Journal of Philosophy* 98, no. 9 (2001): 445–468.
- Hurka, Thomas. "Two Kinds of Satisficing." *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 59, no. 1 (May 1, 1990): 107–111.
- James, William. *The Will to Believe: And Other Essays in Popular Philosophy*. Longmans, Green, and co., 1896.

- Jay Wallace, R. "'Ought', Reasons, and Vice: a Comment on Judith Jarvis Thomson's Normativity." *Philosophical Studies* 154, no. 3 (May 13, 2011): 451–463.
- Jon Elster. "Rationality, Morality, and Collective Action." *Ethics* 96, no. 1 (October 1, 1985): 136–155.
- Joyce, Richard. "Moral Anti-Realism." Edited by Edward N. Zalta. *The Stanford Encyclopedia of Philosophy* (Summer 2009).
<http://plato.stanford.edu/archives/sum2009/entries/moral-anti-realism>.
- Kelly, Thomas. "Epistemic Rationality as Instrumental Rationality: A Critique." *Philosophy and Phenomenological Research* 66, no. 3 (May 1, 2003): 612–640.
- Kelly, Thomas. "Evidence and Normativity: Reply to Leite." *Philosophy and Phenomenological Research* 75, no. 2 (September 1, 2007): 465–474.
- Kelly, Thomas. "Evidence: Fundamental Concepts and the Phenomenal Conception." *Philosophy Compass* 3, no. 5 (2008): 933–955.
- Korcz, Keith Allen. "The Causal-Doxastic Theory of the Basing Relation." *Canadian Journal of Philosophy* 30, no. 4 (December 1, 2000): 525–550.
- Kornblith, Hilary. "Epistemic Normativity." *Synthese* 94 (March 1993): 357–376.
- Kvanvig, Jonathan. "Certain Doubts." *Gibbard Is Not an Expressivist*, September 23, 2010. http://el-prod.baylor.edu/certain_doubts/?p=2205.
- Kvanvig, Jonathan. "Against Pragmatic Encroachment." *Logos & Episteme: An International Journal of Epistemology* (2011).
- Kvanvig, Jonathan. "Coherentism and Justified Inconsistent Beliefs: A Solution." *Southern Journal of Philosophy* (2012).
- Kvanvig, Jonathan L. *The Intellectual Virtues and the Life of the Mind: On the Place of the Virtues in Epistemology*. Studies in Epistemology and Cognitive Theory. Savage, Md: Rowman & Littlefield Publishers, 1992.
- Kvanvig, Jonathan L. "Norms of Assertion." In *Assertion: New Philosophical Essays*. Oxford: Oxford Univ Pr, 2011.
- Kvanvig, Jonathan L. "Millar on the Value of Knowledge." *Aristotelian Society: Supplementary Volume* Supp, no. 85 (January 1, 2011): 83–99.
- Kvanvig, Jonathan L. "A Critique of Van Fraassen's Voluntaristic Epistemology." *Synthese* 98, no. 2 (February 1, 1994): 325–348.

- Kvanvig, Jonathan L. "Propositionalism and the Perspectival Character of Justification." *American Philosophical Quarterly* 40, no. 1 (January 1, 2003): 3–17.
- Lehrer, Keith. *Theory Of Knowledge: Second Edition*. Westview Press, 2000.
- Lehrer, Keith, and Stewart Cohen. "Justification, Truth, and Coherence." *Synthese* 55, no. 2 (May 1, 1983): 191–207.
- Lehrer, Keith, and Thomas Paxson. "Knowledge: Undefeated Justified True Belief." *Journal of Philosophy* 66, no. 8 (1969): 225–237.
- Mackie, J. L. *Ethics: Inventing Right and Wrong*. Penguin (Non-Classics), 1991.
- Mackie, John Leslie. *The Miracle of Theism: Arguments for and Against the Existence of God*. Clarendon Press, 1982.
- Madigan, Timothy J. *W. K. Clifford and the Ethics of Belief*. Cambridge Scholars Publishing, 2008.
- Millar, Alan. "Why Knowledge Matters." *Aristotelian Society: Supplementary Volume* Supp, no. 85 (January 1, 2011): 63–81.
- Montmarquet, James A. *Epistemic Virtue and Doxastic Responsibility*. Rowman & Littlefield Publishers, 1993.
- Mulgan, Tim. "Slote's Satisficing Consequentialism." *Ratio* 6, no. 2 (2006): 121 – 134.
- Pappas, George Sotiros. *Justification and Knowledge: New Studies in Epistemology*. Springer, 1979.
- Plantinga, Alvin. *Warrant and Proper Function*. Oxford University Press, USA, 1993.
- Pollock, John L., and Joseph Cruz. *Contemporary Theories of Knowledge*. Second ed. Rowman & Littlefield Publishers, 1999.
- Radford, Colin. "Knowledge---by Examples." *Analysis* 27, no. 1 (1966): 1–11.
- Raz, Joseph. "Explaining Normativity: On Rationality and the Justification of Reason." *Ratio* 12, no. 4 (December 1, 1999): 354–379.
- Raz, Joseph. "The Myth of Instrumental Rationality." *Journal of Ethics and Social Philosophy* 1, no. 1 (April 1, 2005): 1–28.
- Reed, Baron. "A Defense of Stable Invariantism." *Nous* 44, no. 2 (June 2010): 224–244.

- Riggs, Wayne. "Open-mindedness." *Metaphilosophy* 41, no. 1-2 (January 1, 2010): 172–188.
- Roberts, Robert C., and W. Jay Wood. *Intellectual Virtues: An Essay in Regulative Epistemology*. Oxford University Press, USA, 2007.
- Sartwell, Crispin. "Knowledge Is Merely True Belief." *American Philosophical Quarterly* 28, no. 2 (April 1, 1991): 157–165.
- Schroeder, Mark. "The Scope of Instrumental Reason." *Philosophical Perspectives* 18, no. 1 (December 1, 2004): 337–364.
- Sen, Amartya K. "Rationality and Morality: A Reply to Kurt Baier." *Erkenntnis: An International Journal of Analytic Philosophy* 11 (1977): 225–232.
- Slote, Michael, and Philip Pettit. "Satisficing Consequentialism." *Proceedings of the Aristotelian Society, Supplementary Volumes* 58 (1984): 139–163+165–176.
- Slote, Michael A. *Beyond Optimizing: A Study of Rational Choice*. Harvard University Press, 1989.
- Sobel, David. "The Impotence of the Demandingness Objection." *Philosophers' Imprint* 7, no. 8 (2007): 1–17.
- Sosa, Ernest. *Virtue Epistemology*. Oxford University Press, 2007.
- Sowden, Lanning. "The Inadequacy of Bayesian Decision Theory." *Philosophical Studies* 45, no. 3 (1984): 293 – 313.
- Stalnaker, Robert C. *Inquiry*. A Bradford Book, 1987.
- Thomson, Judith Jarvis. *Normativity*. Open Court, 2008.
- Thorton, Allison. "Knowing and Double Checking." Unpublished Manuscript, 2012.
- Tiberius, Valerie, and Jason Swartwood. "Wisdom Revisited: a Case Study in Normative Theorizing." *Philosophical Explorations* 14 (September 2011): 277–295.
- Turri, John. "On the Relationship Between Propositional and Doxastic Justification." *Philosophy and Phenomenological Research* 80, no. 2 (March 1, 2010): 312–326.
- Tversky, Amos, and Daniel Kahneman. "Judgment Under Uncertainty: Heuristics and Biases." *Science* 185, no. 4157 (1974): 1124 –1131.

Vahid, Hamid. "Truth and the Aim of Epistemic Justification." *Teorema* XXII, no. 3 (2003): 83–91.

Wedgwood, Ralph. *The Nature of Normativity*. Oxford University Press, USA, 2007.

White, Roger. "Epistemic Permissiveness." *Nous-Supplement: Philosophical Perspectives* 19 (January 1, 2005): 445–459.

Williamson, Timothy. *Knowledge and Its Limits*. Oxford University Press, USA, 2002.

Zagzebski, Linda Trinkaus. *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*. Cambridge University Press, 1996.