

A Dutch Book Theorem for Quantificational Credences*
Benjamin Lennertz
Western Kentucky University

In this paper I present an argument for a rational norm involving quantificational credences. To support this norm, I prove a result called a Dutch Book Theorem. In order to prove the result, I introduce the novel concept of a quantificational bet. I also undertake a discussion of Dutch Book Theorems in general and remark on the similarities and differences between the Dutch Book Theorem for quantificational credences and Dutch Book Theorems for norms on ordinary and conditional credences. Overall, the discussion of the norm on quantificational credences gives us a fuller picture of the normative landscape of credal states.

1. Introduction

This paper's focus is a rational norm involving a kind of credal attitude called a quantificational credence¹—the kind of attitude we can report by saying that Lucy thinks that each record in Schroeder's collection is 5% likely to be scratched. I prove a result called a Dutch Book Theorem, which constitutes conditional support for the norm. Though Dutch Book Theorems exist for norms on ordinary and conditional credences, there is controversy about the epistemic significance of these results. So, my conclusion is that if Dutch Book Theorems do, in general, support norms on credal states, then we have support for the suggested norm on quantificational credences. Providing conditional support for this norm gives us a fuller picture of the normative landscape of credal states.

There is another important upshot of my result. It gives us a better picture of the ways in which quantificational credences are similar to and distinct from widely accepted credal attitudes—ordinary and conditional credences. On one hand, featuring in a Dutch Book Theorem is an interesting point in common between quantificational credences and these more standard credal attitudes. On the other hand, a kind of bet at play in proving the Dutch Book Theorem involving quantificational credences—what I call a quantificational bet—is distinct from ordinary and conditional bets (or combinations thereof), which occupy the literature on Dutch Books. Juxtaposing these points develops our picture of the relationship between quantificational credences and ordinary and conditional ones.

The paper proceeds as follows. First, I'll provide a primer on quantificational credences. Then, I'll sketch some background on Dutch Books. Next, I'll introduce the notion of quantificational bets. Finally, I will prove the Dutch Book Theorem for the norm on quantificational credences.

2. Quantificational Credences

In this section I will briefly introduce quantificational credences. Then, I will present a purported rational norm involving quantificational credences.

2.1. What are Quantificational Credences?

Snoopy thinks that each human is somewhat likely to adopt him. The attitude attributed to Snoopy is not an ordinary credence. In particular, it is not a credence of moderate strength in the proposition that each human adopts Snoopy. Having this ordinary credence is thinking it

* Final version to appear in *Ergo*.

¹ Quantificational credences were first discussed in Lennertz (2015).

somewhat likely that everyone adopts him. But surely the rational Snoopy knows that he will only be adopted by one human or family. He thinks it not at all likely that everyone adopts him. So, Snoopy does not have a positive strength credence that everyone adopts him, though he does think that each human is somewhat likely to adopt him. Therefore, our target thought cannot simply be this ordinary credence.

It's also not true that Snoopy thinks that each human is somewhat likely to adopt him simply in virtue of some number of ordinary credences that Snoopy has. In particular, we can't say that Snoopy thinks this just in case, for each object that turns out to be a human, *x*, Snoopy has a moderate strength credence that *x* adopts him. This is because Snoopy likely has no way of having *de re* thoughts about each individual human. This same consideration tells against another initially plausible suggestion—that we can say that Snoopy thinks this just in case, for each object, *x*, Snoopy has a moderate strength conditional credence that *x* adopts him if *x* is a human. Snoopy is even less likely to possess the ability to have *de re* thoughts about each object. The lesson is that for Snoopy to think that each human is somewhat likely to adopt him is neither for him to have some ordinary or conditional credence nor for some fact about his ordinary or conditional credences to be true. The facts about Snoopy's ordinary and conditional credences don't determine whether or not Snoopy thinks that each human is somewhat likely to adopt him.

Given this, what is it for Snoopy to think that each human is somewhat likely to adopt him? It is for him to have what we can call a quantificational credence. A quantificational credence is like a commitment to form some ordinary credence or credences on coming to gain information. For instance, it is the sort of state that commits Snoopy, upon learning that Schroeder is a human, to forming an ordinary credence of moderate strength that Schroeder is going to adopt him. This applies similarly to other objects that he learns are humans. I'll further discuss this characteristic role of quantificational credences in the next subsection, where I address the norms that they place people under. First, I want to say a bit about what typically causes people to form quantificational credences.²

Lennertz (in press) says that the characteristic feature of the functional role of quantificational credences is that they “give us the ability to properly respond to nonspecific probabilistic evidence.” Whereas evidence about a particular human, like Charlie, adopting him would lead Snoopy to form an ordinary credence—like a credence of some strength that Charlie will adopt him—evidence merely about objects of a certain type will lead Snoopy to have a quantificational credence—like his thought that each human is somewhat likely to adopt him. Here are some examples of the sort of non-specific evidence that could lead to forming such a credence. Snoopy might have statistical evidence that suggests that, for each human, there is some chance that that human adopts him. Or he might have general evidence about humans finding him cute and wanting to take him home. Or he may have been informed by the reliable Woodstock that each human is somewhat likely to adopt him. In each of these cases, his evidence is non-specific in that it isn't *de re* evidence about any particular human adopting him. Because of this, Snoopy comes to have a quantificational credence that isn't about any particular object, though it commits him to possible future positions about many objects—including ones he might not at present even know exist.

Before moving on to a discussion of norms, I want to briefly explain the structure of these attitudes. Using the term “quantificational credence” suggests that the attitude under discussion is a partial strength belief with an ordinary content, rather than a full belief with a content that talks about partiality. When thinking about ordinary credences, the distinction can be

2 Thanks to an anonymous reviewer for raising and discussing the issues in the two preceding paragraphs.

exemplified as follows: thinking it 50% likely that the coin lands heads could be a partial belief of 50% strength in the proposition that the coin lands heads or it could be a full belief in the proposition that the probability that the coin lands heads is 50%. I'm not going to argue decisively for the first view here, but I will suggest that a reason that supports thinking that there are genuinely partial ordinary and conditional credal attitudes supports thinking that there are genuinely partial quantificational credal attitudes. That reason goes as follows: it seems that someone without the concept of probability could have these sorts of likely-thoughts. For instance, a very young child—one who lacked the concept of probability—could think it pretty likely that this green mush is going to be yucky or could think it very likely that this green mush is going to be yucky given that it is cold; that is, she could have an ordinary or conditional credence without having the ability to have beliefs about probability. Some find this reason to speak strongly in favor of thinking that likely-thoughts are (at least in many cases) genuinely partial beliefs, or credences. The important point for us is that this reason is just as strong when we consider quantificational likely-thoughts. For instance, the same child could think every lump of green mush is pretty likely to be yucky; that is, she could have a quantificational credence without having the ability to have a belief about probability. I don't take this reasoning to decisively show that it is impossible to make sense of quantificational likely-thoughts as beliefs about quantificational probabilistic propositions. But it does show that if we think that ordinary and conditional credences should be understood as genuinely partial beliefs, there is very good reason to think that quantificational credences should be as well.³

In this subsection, I've introduced a way of thinking about what quantificational credences are. We can further improve our understanding of these attitudes by seeing the sort of rational role that they play. So, the ensuing discussion of norms is meant to both introduce the subject for my later arguments and help us better understand quantificational credences and the role they play in the mental life of the uncertain agent.

2.2. A Norm on Quantificational Credences

Suppose, again, that Snoopy has the following quantificational credence: he thinks that each human is somewhat likely to adopt him. Suppose that he also thinks that some human is not at all likely to adopt him. It seems that such a combination is irrational. The following is a slightly different irrational combination. Snoopy has the original quantificational credence, he knows that Sally is a human, and he has a credence of strength 0 that Sally will adopt him. Again, it seems like such a combination is irrational.

We can provide a norm to explain these instances of irrationality. In order to understand it, we need to introduce a more technical way of talking about quantificational credences. We can call Snoopy's thought that each human is somewhat likely to adopt him *a quantificational credence of some strength, for each human, that that human adopts him*. In general, for some quantifier, q , properties, F and G , and real number, n , we can call a thought that q of the F s are n -likely to be G s *a quantificational credence of strength n , for q of the F s, that they are G s*.

We will build up to the norm by introducing a few new notions.⁴ First, we need the notion of a *totally decided set of beliefs*—a set of beliefs that includes, for every proposition, a belief in either that proposition or its negation—and a *totally decided set of credences*—a set of credences that includes, for every proposition, an ordinary credence in that proposition. Next, we need the notion of satisfying a quantificational credence. We can say that a pair consisting of a totally

3 Thanks to an anonymous reviewer for urging me to confront the issues raised in this paragraph.

4 This construction is a simplified version of that carried out in Lennertz (2015).

decided set of beliefs, B , and a totally decided set of credences, C , *satisfies* a quantificational credence of strength n , for q of the F s, that it is G if and only if for each q of the objects believed to be F in B , x , C contains a credence of strength n that x is G . For instance, a pair consisting of a totally decided set of beliefs and a totally decided set of credences satisfies a thought that each human is somewhat likely to adopt Snoopy if and only if for each object believed to be a human, there is a credence of some strength that that object adopts Snoopy. Finally, we need the notion of a complete extension of a set of beliefs and a set of credences. The general idea here is to extend an actual set of beliefs or credences to be totally decided in a way that does not introduce any irrationality. More precisely, a *complete extension of a set of beliefs* is a superset of those beliefs that is totally decided and that includes no further beliefs about those propositions in the original set or their negations. A *complete extension of a set of credences* is a superset of those credences that is totally decided and that includes no further credences about those propositions or their negations.

We can now state the norm. Informally, the idea is that rationality requires a possible way of coming to have new beliefs and credences that allow one to satisfy one's quantificational credences. So, one is irrational if one can't extend one's beliefs and credences in ways that allow one to satisfy one's quantificational credences. More precisely:

Quantificational Consistency: An agent, A , should be such that there is a pair consisting of a complete extension of A 's beliefs and a complete extension of A 's credences that satisfies A 's quantificational credences.^{5,6}

We can better understand this norm by seeing how it applies to the two examples sketched above. In the first, Snoopy thinks that each human is somewhat likely to adopt him and thinks that some human is not at all likely to adopt him. Snoopy is irrational in virtue of these two quantificational credences. Quantificational Consistency predicts this. We've made no special assumptions about Snoopy's beliefs or ordinary credences; the quantificational credences are inconsistent even without any such assumptions. This is because there is no pair consisting of a totally decided set of beliefs, B , and totally decided set of credences, C , that satisfies both of these quantificational credences. In order to satisfy the first, it must be that, for all objects, x , such

5 The norm that I presented in Lennertz (2015) is more complex because I incorporated conditional credences. This is to account for the fact that quantificational credences can clash with conditional ones. I have left them out to reduce the complexity of our discussion. But we could carry the result from Section 5 over to this more complex norm.

6 Quantificational Consistency mentions both qualitative beliefs and credences. The way the notion of satisfaction functions in the norm requires that beliefs and credences exist and are independent of each other (in the sense that there is a possible pair of beliefs and credences that contains a belief in p and a very low credence in p). Many theorists want to reduce beliefs to credences which have some sort of property—perhaps credences whose strengths fall above some threshold. Could such theorists go in for something like Quantificational Consistency? Yes, and though the details are complicated, I'll give a quick sketch. The new norm would have to be more complex than Quantificational Consistency, since we would no longer be able to independently adjust beliefs and credences when checking for possible extensions that satisfy a quantificational credence. Thus, we might miss out on belief-credence pairs that could satisfy a quantificational credence. We can skirt the problem by complicating our procedure. We would need a notion of satisfaction that didn't require complete extensions of beliefs and credences. Instead of looking at a single complete extension of all of an agent's beliefs and credences, we would look at multiple smaller sets of beliefs and credences to avoid any places where the proposed belief-credence connection would constrain our search in the way suggested above. In general, we would have to be more selective about what attitudes we hold fixed in each extension, and then we would have to perform multiple extensions, piecing the results together in the appropriate manner. A formal presentation of this sort of picture would take us too far afield from our present course. Thanks to an anonymous referee for raising this issue.

that the proposition that x is a human is in B, C contains a moderate strength credence that x adopts Snoopy. In order to satisfy the second, it must be that, for some object, y , such that the proposition that y is human is in B, C contains a credence of 0 strength that y adopts Snoopy. There are no sets B and C that can meet both of these demands. Thus, Snoopy's quantificational credences violate Quantificational Consistency; he is irrational.

The other case works in a somewhat different way. This was the case where Snoopy thinks that each human is somewhat likely to adopt him, knows that Sally is a human, and has a credence of 0 strength that Sally will adopt him. So, any complete extension, B , of Snoopy's set of beliefs will contain the belief that Sally is human and any complete extension, C , of Snoopy's set of credences will contain a credence of 0 that Sally will adopt him. So, it is then impossible for $\langle B, C \rangle$ to satisfy his quantificational credence. That is, it cannot be that, for all objects, x , such that the proposition that x is a human is in B, C contains a moderate strength credence that x adopts Snoopy. The combination of his belief and credence about Sally make this impossible. So, he violates Quantificational Consistency and is irrational.

Now that we have seen some examples involving Quantificational Consistency, we are in position to offer considerations in its favor. In order to do so, I will first outline a kind of result that is used to support norms on credal states.

3. Dutch Book Theorems

Dutch Book Theorems are used to support norms of rationality involving credal attitudes. They are results that show that any agent who has a credal state (or succession of credal states) that violates a given norm is committed to making bets that are guaranteed to lose her money. Many philosophers make the further step of claiming that the existence of a Dutch Book Theorem for a given norm supports the norm as a principle of rationality.⁷ In this section, I'll introduce Dutch Books and their relationship to rational norms.

3.1. Credences and Fair Betting Prices

Before giving an example to clarify the idea of Dutch Books, I will discuss the connection between one's credences and fair betting prices. The assumed connection is that one's credences *sanction as fair* particular bets.⁸ A credence of strength n in p sanctions as fair a bet where one pays $\$n$ and wins $\$1$ when p is true and $\$0$ when p is false. This credence also sanctions as fair what we might call *selling* the same bet: taking $\$n$ to pay out $\$1$ when p is true and $\$0$ when p is false. For instance, a credence of .6 that it is sunny in Shanghai sanctions as fair a bet of $\$.6$ to win $\$1$ if it is sunny in Shanghai and $\$0$ if it is not sunny in Shanghai. That credence also sanctions as fair selling the same bet. Furthermore, that credence sanctions as fair scaled versions of these bets. For instance, we could scale the above bets by 10 to get a bet of $\$6$ to win $\$10$ if it is sunny in Shanghai and $\$0$ if it is not.

If we think about the sorts of bets a rational agent might actually accept, the above

7 Ramsey merely mentions, but does not develop, this line of thought: "Having degrees of belief obeying the laws of probability implies a further measure of consistency, namely such a consistency between the odds acceptable on different propositions as shall prevent a book being made against you" (1964: 81). De Finetti (1964: 101–103) endorses this sort of argument.

8 I follow Christensen (1996) in using the notion of a credence sanctioning a bet as fair. I agree with him that there is a natural sense in which this is so (even if a particular agent may not avow the fairness of particular bets). Hedden (2013) assimilates the licensing of bets by an agent's credences to the licensing of action in general by those credences in a decision theoretic framework. Though such an account has the theoretical virtue of offering a unified picture of bets and other action, it, as Hedden shows, allows that there are agents with incoherent credences who are not susceptible to Dutch Books.

discussion idealizes in at least three ways. First, it assumes money to have linear value for an agent. But it may be rational to value gains and losses of different amounts of money differently (Hájek 2008b). Second, it assumes that bets with an expected utility of 0 are sanctioned as fair. But it might be equally rational to for an agent to be indifferent to such bets and refuse them, since there is no expected gain (Hellman 1997: 195; Vineberg 2016). Third, it assumes that agents are not rationally risk averse. But perhaps it is rational to refuse bets with positive expected value where the possible losses are, though unlikely, very great (Buchak 2013; 2014). Others have discussed these worries about the setup of Dutch Books, and I will ignore them in what follows. As I've said, my goal is to argue that, insofar as Dutch Book Arguments are successful, we can support Quantificational Consistency using one.

3.2. An Example: Additivity and Dutch Books

In this subsection, I'm going to give an example of a Dutch Book for an agent with credences that violate a (purported) norm on rationality. That is, I'll give a series of bets which the agent's credences sanction as fair, but which jointly guarantee the agent a loss no matter how the world turns out to be.⁹ Our example will focus on the following norm on an agent's credal state:

Additivity: The strength of an agent's credence in the disjunction of two inconsistent propositions should be equal to the sum of the strength of her credence in the first and the strength of her credence in the second.

This norm corresponds to a restricted version of the additivity axiom of probability theory. There is a Dutch Book Theorem for this norm:

Dutch Book Theorem for Additivity: If an agent's credal state violates Additivity, then there exists a Dutch Book against her.

Let's consider an example where an agent violates this norm and is susceptible to a Dutch Book.

Suppose that Charlie has received a single valentine and does not know who it is from. Suppose that he has a credence of strength .3 that the valentine came from Patty and a credence of strength .2 that the valentine came from Marcie. Further suppose that Charlie has a credence of strength .4 that the valentine came from Patti or that the valentine came from Marcie. This combination of credences violates Additivity, since the strength of his credence in the disjunction (of two inconsistent propositions) is less than the sum of the strengths of his credences in the disjuncts.

We can arrange a series of bets, each of which is sanctioned as fair by one of his credences, that guarantees a sure loss to Charlie. First, we have him bet \$3 that the valentine is from Patty. If it is from her, he wins \$10 and if it is not from her, then he wins \$0. Charlie's credence of strength .3 that the valentine came from Patty sanctions this bet as fair. Similar considerations hold for the other bets. For the second bet, we have him bet \$2 that the valentine is from Marcie. Again, if it is from her, he wins \$10 and if it is not from her, then he wins \$0. Finally, we pay \$4 to Charlie in a bet that will pay us \$10 if the valentine is from Patty or the valentine is from Marcie and \$0 if the valentine is from neither. This series of bets guarantees Charlie a loss. If the valentine is from Patty, Charlie gains \$7 from the first bet but loses \$2 from the second and \$6 from the third. This amounts to a loss of \$1. If the valentine is from Marcie,

9 Mahtani (2015) gives an alternative way of thinking about the guaranteed loss—one that looks at whether the agent will lose money given every possible interpretation of the non-logical symbols in the contents of the credences. There are interesting differences between this view and the no-matter-how-the-world-turns-out view, both in general and with regard to the topics in this paper. But for reasons of space, I must save such comparisons for another place.

Charlie gains \$8 from the second bet but loses \$3 from the first and \$6 from the third—again, a loss of \$1. Finally, if the valentine is from neither, Charlie makes \$4 from the third bet but loses \$3 from the first and \$2 from the second. So, however the world turns out, Charlie is guaranteed a loss of \$1. That is, he is susceptible to a Dutch Book.

This sequence of bets against Charlie constitutes a Dutch Book. But the fact that Charlie is susceptible to a Dutch Book is not dependent on any of the particular features of our case. In fact, it can be shown that anyone who violates Additivity is susceptible to a Dutch Book—i.e., that the Dutch Book Theorem for Additivity is true. Similar results have been shown for a number of other purported norms on credal states.¹⁰ I won't rehearse the proof of the Dutch Book Theorem for Additivity or these other proofs here, though I will prove a general result for Quantificational Consistency.

3.3. What do Dutch Book Theorems show?

Alan Hájek (2005) is explicit about the distinction between Dutch Book Theorems and Dutch Book Arguments. A Dutch Book Theorem is a result that says that if an agent has a credal state with some particular property, there exists a Dutch Book for that agent. The Dutch Book Theorem for Additivity, mentioned above, is an example. A Dutch Book Argument takes this result and claims that it shows that it is irrational to have a credal state with that property.

Many who have accepted the existence of Dutch Book Theorems have doubted the force of Dutch Book Arguments. Suppose we accept that if an agent has a credal state that has some property, then there is a series of bets that her credences sanction as fair that guarantee her a loss (as a Dutch Book Theorem would say). Still, it does not follow that an actual agent with that credal state would be subject to sure losses. This would require there actually being nefarious bookies around (Christensen 1991: 237; 1996: 451; Vineberg 1997: 186). In addition, it would require the agent to not pretend that she has other credences in betting in order to avoid a loss or gain other benefits (Christensen 1991: 237; Vineberg 1997: 186; Hájek 2008a calls this masking her dispositions; Hájek 2008b). It would also require that an agent does not get wise to such a bookie and change her credences in the middle of the series of bets being offered (thus resolving the previous inconsistency).¹¹

However, the most important worry about the rational import of Dutch Book Theorems is that there seems to be a gap between facts about a loss in a series of bets and epistemic irrationality. As David Christensen says, “even if probabilistically incoherent agents were subject to real practical difficulties, it would not obviously follow that their beliefs were defective from the *epistemic* standpoint” (1996: 451). Furthermore, some doubt that it could show much about rationality at all, since just as there are fair bets that the agent will accept that will guarantee her a loss, there are fair bets that the agent will accept that will guarantee her a gain.¹²

A classic response to the main worry is to claim that the guarantee of a practical loss does not constitute the epistemic irrationality of the agent. Rather the guarantee of a practical loss indicates or dramatizes some epistemic inconsistency in the agent's credal state (Christensen 1991; 1996). As Christensen says:

10 These include norms corresponding to all of the probability axioms (Skyrms 1986: 185–188), a diachronic conditionalization norm (Lewis 1999; Teller 1973: 222–225; Skyrms 1986: 189–192), and the principle of reflection (van Fraassen 1984).

11 See Jackson and Pargetter (1976) and Hellman (1997) for attempts to strip away all of the practical elements of Dutch Books, leaving only the epistemic lessons.

12 See Hájek (2005; 2008a) for discussion of this (what he calls Good or Czech Books) and his solution.

It is also intuitively plausible that, if a set of betting odds allows someone to devise a priori a way of exploiting those odds to inflict a sure loss, then there is something amiss with those betting odds. And finally, if a single set of beliefs sanctions as fair each of a set of betting odds, and that set of odds is defective, then there is something amiss with the beliefs themselves. The fact that the diagnosis can be made a priori indicates that the defect is not one of fitting the beliefs with the way the world happens to be. Like deductive inconsistency, it is a defect internal to the agent's belief system.¹³ (1996: 457)

I'm inclined to think that Dutch Book Theorems do tell us about rational norms, but I won't try to defend this position in this paper. Instead, I will settle for a conditional conclusion: If Dutch Book Theorems do support rational norms, then there is support for our rational norm on quantificational credences.¹⁴ This is because, as I will show, there is a Dutch Book Theorem for Quantificational Consistency.

4. Quantificational Bets

In order to be able to form our Dutch Book Theorem for Quantificational Consistency we will need to use a heretofore undiscussed sort of bet—a quantificational bet. This is the sort of bet that is sanctioned as fair by a quantificational credence.

Suppose that Lucy has a quantificational credence of strength .05, for each record in Schroeder's collection, that it is scratched. This attitude sanctions as fair the following quantificational bet (or, at least, something very much like it—the details of this qualification will be discussed in two paragraphs). Given a finite domain of objects, we have Lucy pay \$5 for each object in the domain. Then, when Lucy comes to have beliefs about which objects are all and only the records in Schroeder's collection, we proceed as follows. If a given object is not believed to be a record in Schroeder's collection, then the \$5 for that object is reimbursed. If a given object is believed to be a record in Schroeder's collection, then Lucy gets the following payout: \$100 if that record is scratched and \$0 if that record is not scratched. If we are worried about such practicalities as Lucy not having enough money upfront, then we can run the same sort of process but, instead of her paying \$5 for each object, we mark down in a ledger that she owes \$5 for each object. We then cross out that debt for each object that she turns out not to believe to be one of the records in Schroeder's collection. For each of the objects that she does come to believe are the records in Schroeder's collection, she must make good on her \$5 payment. From there, the payout structure is the same as the original presentation of the bet.

Lucy's quantificational credence will sanction this bet as fair. This is for the following

13 Christensen (1991: 239) makes similar remarks. However, readers should see Maher (1997) for a criticism of Christensen's position. Vineberg (2001) is worried about similar issues—where we need to account not just for each of the bets being sanctioned as fair individually, but their being sanctioned as fair collectively. Skyrms (1986) makes this assumption of a Dutch Book explicit. And Hellman (1997: 196–7) provides a response for projects in the philosophy of science.

14 The literature is often more nuanced than represented here—where Dutch Books Arguments either succeed in supporting norms or they do not. In particular, many authors make a distinction between Dutch Books that include bets at a single time—synchronic Dutch Books—and Dutch Books that include bets at multiple times—diachronic Dutch Books (sometimes called Dutch Strategies). The evidential import of diachronic Dutch Books is often taken to be more suspect than the import of synchronic ones—one reason being that there are diachronic Dutch Book Theorems for norms that seem obviously false, like the principle of reflection (van Fraassen 1984). See Christensen (1991; 1996), Maher (1992), Hellman (1997), and Mahtani (2012). See Armendt (1992), Vineberg (1997), Briggs (2009), and Mahtani (2015) for positions which claim that the fundamental problem with the Dutch Book Theorem for the principle of reflection is not its diachronicity. In Section 5.2 I'll discuss why my Dutch Book Theorem is not subject to these worries.

reasons: her credence says nothing about objects that are not believed to be one of the records in Schroeder's collection, and the quantificational bet does not, once Lucy's mind is made up, concern them. About each object that she does believe is a record in Schroeder's collection, Lucy is committed by her quantificational credence to thinking that it is 5% likely to be scratched. So, her attitude sanctions as fair a sub-wager on that object, which is just equivalent to a \$5 bet on that object being scratched (to win \$100). So, the quantificational bet is just the commitment to make, once Lucy is decided about which objects are the records in Schroeder's collection, the combination of such bets on these objects.

However, the bet just described isn't exactly the one that Lucy's quantificational credence sanctions as fair. Really, the bet sanctioned as fair is like the one just described but with an important restriction. The bet goes in the way described in a very specific scenario—one where Lucy comes to learn which objects are the records in Schroeder's collection and learns nothing else. In scenarios where Lucy learns other facts, the bet is simply called off and all money is refunded. Why make this restriction? Because cases where Lucy learns other facts are cases where she may be justified in revising her quantificational credence. In originally forming that attitude, Lucy isn't committing to sticking with it come what may. She may get new evidence that some of the records in Schroeder's collection definitely aren't scratched. So, she may revise her quantificational credence to be one of strength .05, for most (rather than all) records in Schroeder's collection, that they are scratched. Since Lucy's credal state can rationally evolve in this way, the bet that her quantificational credence sanctions as fair is one that only continues to have grip if what she learns is not evidence against her quantificational credence. The easiest way to formalize this is to say that it remains in force in cases where all that she learns is what records are all and only those in Schroeder's collection. Though this restriction will be in place throughout the discussion of quantificational bets, I will often omit explicit mention of it except where it is particularly relevant.

Making this quantificational bet is equivalent to making, for each object in the domain, a conditional bet of \$5 (to win \$100) that it is scratched if it is a record in Schroeder's collection. For a given object, o , the conditional bet of \$5 (to win \$100) that it is scratched if it is a record in Schroeder's collection is one which has the following payout structure: Lucy's net payout is \$0 if she learns that o is not a record in Schroeder's collection (above I described this as the \$5 being reimbursed) or if she learns facts beyond those about whether objects are records in Schroeder's collection.¹⁵ Lucy's net payout is \$-5 if she learns only that o is a record in Schroeder's collection but it is not scratched (she puts up \$5 and wins \$0). Lucy's net payout is \$95 if she learns only that o is a record in Schroeder's collection and it is, in fact, scratched (she puts up \$5 and wins \$100). When we consider the collection of bets with this payout structure for each object in the domain, we have something equivalent to the quantificational bet explained above. In general, making the sort of quantificational bet that is sanctioned as fair by a universal quantificational credence is equivalent to making a set of conditional bets.

However, not all quantificational bets are equivalent to sets of conditional bets. In particular, the bet sanctioned as fair by what we can call a less-than-universal quantificational

15 The reasons for making the restriction I discussed in the previous paragraph apply in much the same way to conditional credences. A conditional credence of strength n in p if q sanctions as fair a bet that commits an agent to betting on p at odds determined by n if she learns q . But her commitment does not require her to go through with this wager if she learns many other things as well, some of which may give her reason to revise her original conditional credence. This is why diachronic conditionalization norms are often phrased in terms of what credence an agent should have in p at t_2 if she has a conditional credence of n in p if q at t_1 and comes, at t_2 , to learn only q .

credence is not equivalent to some particular set of conditional bets. Consider the following example: Charlie Brown's baseball team is notoriously bad. Linus knows that it is Charlie's pitching that makes them so bad and that the team sometimes does well when Charlie isn't there. Linus thinks that Charlie will miss exactly three games during the season. So, he thinks that (at least and at most) three games are 50% likely to be won by Charlie Brown's team. This quantificational credence—about three games—has less-than-universal quantificational strength.

This less-than-universal quantificational credence sanctions as fair the following wager. For each object in the domain, we put Linus down for a \$50 commitment in the ledger. Let's suppose that Linus comes to have new beliefs about which objects are all and only those that are baseball games and does not form new beliefs about anything else. If this were a bet sanctioned as fair by a universal quantificational credence, then for every object that Linus comes to believe is one of the baseball games, he would owe \$50. But it is a less-than-universal quantificational credence. So, Linus merely owes \$50 for three of the objects that he comes to believe are baseball games. And, since he wasn't committed to which particular games are 50% likely to be wins, he can choose which games to wager the \$50 on, as long as he does it for exactly three of them (presumably he'll chose the games that Charlie can't attend—though nothing about the attitude itself requires this). For each of the three games for which he chooses to pay \$50, he wins \$100 if Charlie Brown's team wins that game and he wins \$0 if Charlie Brown's team loses that game. This is how a less-than-universal quantificational bet works.

This less-than-universal quantificational bet is not equivalent to any particular set of conditional bets. It is clearly not equivalent to a set of conditional bets about all objects, since Linus is not committed to having money riding on the outcome of each baseball game. Additionally, there is no particular bunch of conditional bets about some subset of objects that does the job. An easy way to see this is that Linus has a choice, once he learns which objects are baseball games, about which three of them he wants to use to make good on his \$50 commitments. There could be no such choice about purely conditional bets. In general, a less-than-universal quantificational bet is not equivalent to some particular set of conditional bets. It is a distinct sort of bet.

Given that they are not, in general, reducible to a set of conditional bets, we might wonder whether there is something illicit or ill-formed about quantificational bets. Rest assured that there is not. They have a well-defined payment and payout structure. One way in which they differ from many bets is that there is a choice point after the initial bet is made at which certain parts of the bet can be turned off and turned on. But this doesn't seem different in kind from what happens in some casino games. For instance, in craps one can turn off and on many parts of one's bet and one can take some parts back entirely (and put the money back in one's pocket). Just like in the description of a quantificational bet, which parts of the bet are open to be made, canceled, and turned on or off in the middle of the game depend on the original bet made. Another way in which quantificational bets might seem different from the sorts of bets that we commonly think of is that they don't, when made, involve betting on particular objects having a property. But, it's easy to imagine a casino game that works like this. Imagine a game where five cards are drawn and laid face down. In this game one can make a commitment to pay \$5 for each black card that comes up (and \$0 for each red card) with the following payout structure: win \$100 if that black card is an ace and win \$0 if that black card is not an ace. But this is just equivalent to what I have been calling a universal quantificational bet. One could also play a game with a similar betting structure where the dealer looks at the cards, offers you the black cards face down, and you are committed to paying \$5 for all but one of them. That is, you must discard one of them, but you get to choose which. We can assume that the payout structure from

this point on is the same as for the above game. Such a bet has the structure of a less-than-universal (all-but-one) quantificational bet. These games—while too clumsy to appear on the floor of the Bellagio—are surely betting games. And the sorts of bets that are made in playing them are quantificational bets. So, quantificational bets are genuine, well-formed bets.

Furthermore, this discussion of quantificational bets applies in the world outside of our artificial casino floor. Just as some sorts of choices made under uncertainty dramatize real life versions of ordinary bets, other sorts of choices do the same with quantificational bets. In the ordinary case, making a choice to go to the party is, in a sense, a bet that performing that action will lead to better results than staying home. And one often makes this choice on the basis of expected outcome calculations that also appear in explicit betting situations. For instance, Franklin may think that the fun of hanging out with his friends weighted by the likelihood that they attend the party outweighs the boredom of going to a party where he doesn't know people weighted by the likelihood that his friends skip the party—and it does so to a degree that is larger than the assured moderate pleasure of staying home and watching Netflix. But not only do some choices to act have the structure of ordinary bets, some choices to act have the structure of quantificational bets. Consider Pig Pen on his first day of school making the choice to try to be friends with each classmate who chats with him. This is a sort of commitment to attempt to foster friendships with (possibly many) people he does not yet know; so, it is not an ordinary choice to perform some particular action—not the sort of choice that is a real life application of an *ordinary* bet. Instead, the choice has the structure of a quantificational bet. We can see this by understanding why Pig Pen may make this choice. He may think that anyone who chats with him is quite likely to be a nice person. And he may highly value being friends with anyone who is nice. Based on that quantificational credence and those utilities, it may be that attempting to be friends with anyone who chats with him has a better expected outcome than all alternative policies. Pig Pen's reasoning here is structurally similar to that of the agents making the quantificational bets that we investigated above. Thus, like their ordinary, non-quantificational counterparts, such bets need not be seen merely as confined to artificial games of chance. They have the same structure as real-life choices made to perform real-life actions.

Now that we have a better grasp on what quantificational bets are, we are in a position to prove the Dutch Book Theorem for Quantificational Consistency.

5. The Dutch Book Theorem for Quantificational Consistency

We can now prove a Dutch Book Theorem for Quantificational Consistency. Remember that Quantificational Consistency says the following:

Quantificational Consistency: An agent, *A*, should be such that there is a pair consisting of a complete extension of *A*'s beliefs and a complete extension of *A*'s credences that satisfies *A*'s quantificational credences.

The theorem, then, looks as follows:

Dutch Book Theorem for Quantificational Consistency: If an agent's credal state violates Quantificational Consistency, then there exists a Dutch Book against her.

I'll start the section with some examples where someone violates Quantificational Consistency and is susceptible to a Dutch Book. Then, I'll present a note about the nature and structure of our Dutch Book. Finally, I'll prove the Dutch Book Theorem for Quantificational Consistency.

5.1. Examples of Dutch Books for Agents Who Violate Quantificational Consistency

Charlie wants to kick the football to the moon, but he realizes that Lucy almost always pulls the ball away the moment before he tries to kick it. So, he has a quantificational credence, a

thought that each attempt is unlikely—for precision let’s assume 5% likely—to be one where he kicks the football to the moon. But suppose he is, at some particular time, overcome by optimism that this attempt, call it Attempt X, will be one where he kicks the football to the moon. So, he also has an ordinary credence, a thought that it is very likely—for precision, let’s assume 99% likely—that Attempt X is one where he kicks the football to the moon. Furthermore, Charlie believes that Attempt X is an attempt. Charlie is obviously irrational. He violates Quantificational Consistency.

We can form a Dutch Book, showing that Charlie’s credences and quantificational credences sanction as fair a series of bets that guarantee Charlie a loss. First we make a bet that is sanctioned as fair by his ordinary credence. We have Charlie pay us \$99 to win \$100 if Attempt X is one where he kicks the football to the moon and to win \$0 if Attempt X is not one where he kicks the football to the moon.

Next we use a quantificational bet in our Dutch Book. We buy from Charlie a \$5 quantificational bet that, for each attempt, it is one where he kicks the football to the moon. That is, we mark down in the ledger that we owe Charlie \$5 for each object in the domain that he knows is an attempt. He pays us \$100 for each attempt in which he kicks the football to the moon, and he pays us \$0 for each attempt in which he does not kick the football to the moon. As discussed in Section 4, a quantificational credence like Charlie’s will sanction this bet as fair.

We can’t simply combine this quantificational bet with the bet sanctioned by his ordinary credence. The result would not form a Dutch book. It is true that one of the ledger entries for the quantificational bet is equivalent to a conditional bet of \$5 that Attempt X is one where he kicks the football to the moon, given that Attempt X is an attempt. That conditional sub-wager (of the quantificational bet), together with the ordinary bet just mentioned, does (given Charlie’s learning only that Attempt X is an attempt) guarantee a loss with respect to these bets on Attempt X, no matter how the world turns out. We can see this as follows: If Attempt X is one where Charlie kicks the football to the moon, then he wins \$1 from the ordinary bet and loses \$95 from the conditional sub-bet, for a total \$94 loss. If Attempt X is not one where Charlie kicks the football to the moon, then he loses \$99 from the ordinary bet and wins \$5 from the conditional sub-bet, for a total \$94 loss.

But, the quantificational bet comprises conditional sub-bets on every object in the domain. If there are other objects in the domain that Charlie learns to be attempts, then these bets won’t simply be called off; they will pay off one way or another. And Charlie might win back all of his money and more through them. So, Charlie would not be guaranteed a loss—as is necessary for a Dutch Book. Because of this, we can’t simply combine the quantificational bet and the ordinary bet to form a Dutch Book. But, if we were able to isolate that one part of the quantificational bet and combine it with the ordinary bet, we would be able to Dutch Book Charlie in the way we saw in the previous paragraph.

Indeed, we can isolate it by introducing a quantificational bet that cancels out everything except the required conditional sub-bet. We just need to have Charlie make something like the reverse quantificational bet. It will be a \$5 quantificational bet that, for every attempt except Attempt X, it is one where he kicks the football to the moon. And instead of buying the bet from Charlie, we sell the bet to him. That is, we mark in the ledger that he owes us \$5 for each object in the domain that is an attempt other than Attempt X. For each such object, we pay him \$100 if it is an attempt where he kicks the football to the moon and we pay him \$0 if it is an attempt where he does not kick the football to the moon.

Charlie’s quantificational credence sanctions this bet as fair. To see why, we need to understand that a universal quantificational credence sanctions as fair not only quantificational

bets but also a great many conditional bets. Suppose that I have a universal quantificational credence of strength n , for each F , that it is G . This sanctions as fair a conditional bet of $\$n$ on Go given Fo . That is, it sanctions as fair either buying or selling a bet of $\$n$ that pays $\$1$ if Fo and Go are true, pays $\$0$ if Fo is true and Go is false, and is called off (and the $\$n$ refunded) if Fo is false. Why does it sanction this as fair? Because if I have that quantificational credence and I learn that Fo is true, I am committed to thinking that Go is n -likely to be true. But there is nothing special about the choice of the object, o . Such reasoning holds for every object. So, my quantificational credence sanctions as fair a conditional bet of this sort for every object. This reasoning applies to the quantificational bet mentioned in the previous paragraph. Charlie's quantificational credence sanctions as fair buying and selling conditional bets on every object, where $\$5$ is paid to win $\$100$ if that object is attempt where he kicks the football to the moon and win $\$0$ if it is an attempt where he does not kick the football to the moon (and the $\$5$ is refunded if it is not an attempt). Grouping together all of these conditional bets except the one about Attempt X yields the desired reverse quantificational bet. So, it is sanctioned as fair by Charlie's quantificational credence.¹⁶

Importantly, Charlie's reverse quantificational bet balances out the first quantificational bet with respect to every object other than Attempt X. For every attempt other than Attempt X, we both give $\$5$ to and take $\$5$ from Charlie (for a net exchange of $\$0$). If that attempt is one where he kicks the ball to the moon, we both win $\$100$ from and lose $\$100$ to Charlie (for a net exchange of $\$0$). Likewise, if it is one where he does not kick the ball to the moon. What is left is equivalent to buying a conditional bet from Charlie for $\$5$ that Attempt X is one where he kicks the football to the moon, given that it is an attempt. This, together with Charlie's ordinary bet of $\$99$ that Attempt X is one where he kicks the football to the moon, guarantees him a loss, as we saw four paragraphs ago. So, if Dutch Book Arguments show what they purport to, then we have shown that Charlie is irrational.

There is a sense in which this procedure, involving quantificational bets that cancel out parts of other quantificational bets, is unnecessarily complex. We could reach the same result by simply using the fact that we just saw—that a universal quantificational credence sanctions as fair each conditional sub-bet of the quantificational bet sanctioned as fair by that attitude. The canceling of large parts of the quantificational bet allows us to mimic a single conditional bet. And this is enough to form the Dutch Book. However, using this fuller method allows us to assimilate cases that include only universal quantificational credences to those that include less-than-universal quantificational credences. As the following example will show, cases involving the latter sort of attitude require the more complicated strategy.

The example will also be different from the one just given along another dimension—one which was discussed in Section 2.2. The case just given was one where a quantificational credence is inconsistent with some ordinary credence, given the agent's beliefs. The one I'm about to give involves an inconsistency just in virtue of quantificational credences. Nonetheless, it is the existence of the less-than-universal quantificational credence (and not this latter difference) that calls for the more complex strategy.

Consider again Charlie Brown's terrible baseball team. Suppose that, because the team is so awful, Linus thinks that each baseball game is 1% likely to be won by Charlie Brown's team. But suppose, again, that Linus also knows that the team sometimes does well when Charlie can't

¹⁶ I ignore very general worries about whether each of these conditional bets being sanctioned as fair means that all, or all but one, are sanctioned collectively, since these are not unique to the issue being discussed here, but apply to all Dutch Books. See Footnote 13 for discussion.

play. And he thinks that Charlie will miss three games. So, he thinks that three games are 50% likely to be won by Charlie Brown's team. In this situation, Linus has a credal state that violates Quantificational Consistency. There is no extension of his beliefs and credences that satisfies his quantificational credences. We can show that he is irrational by subjecting him to a Dutch Book.

Linus's first quantificational credence sanctions as fair a quantificational bet of \$1, for each baseball game, that it is won by Charlie Brown's Team. In the ledger, we mark down our commitment to paying Linus \$1 for each baseball game (for each object in the domain that we come to know is not one of the baseball games, we owe nothing). Linus pays us \$100 for each baseball game that is won by Charlie Brown's team, and he pays us \$0 for each baseball game that is not won by Charlie Brown's team.

We then offer Linus the following quantificational bet, which is sanctioned as fair by his second quantificational credence. We provisionally put him down in the ledger for \$50 for each baseball game. But the provision is that he must cancel the wager for all-but-three of the objects he learns are the baseball games. This is because he only has a commitment regarding three such games. So, there will remain, once he learns what the baseball games are, three for which he owes \$50. It is up to Linus which three will remain part of the wager. For each of these, he wins \$100 if it is a game that is won by Charlie Brown's team and wins \$0 if it isn't.

Given these two quantificational bets and the fact that Linus will learn that some objects are baseball games, he is committed, for at least one baseball game (three, actually), to actions that are equivalent to betting, at two different odds, on it being a win for Charlie Brown's team.¹⁷ Consider the following charts, which display Linus's betting commitments in light of his quantificational bet after learning that, say, objects a-f are all and only the baseball games:

Linus's universal quantificational credence commits him to all of the bets in this chart:

Object	Payout to Linus	
a	\$-99 if Charlie Brown's team wins a	\$1 if Charlie Brown's team loses a
b	\$-99 if Charlie Brown's team wins b	\$1 if Charlie Brown's team loses b
c	\$-99 if Charlie Brown's team wins c	\$1 if Charlie Brown's team loses c
d	\$-99 if Charlie Brown's team wins d	\$1 if Charlie Brown's team loses d
e	\$-99 if Charlie Brown's team wins e	\$1 if Charlie Brown's team loses e
f	\$-99 if Charlie Brown's team wins f	\$1 if Charlie Brown's team loses f

Linus's less-than-universal quantificational credence commits him to exactly three of the bets in this chart:

Object	Payout to Linus	
a	\$50 if Charlie Brown's team wins a	\$-50 if Charlie Brown's team loses a
b	\$50 if Charlie Brown's team wins b	\$-50 if Charlie Brown's team loses b
c	\$50 if Charlie Brown's team wins c	\$-50 if Charlie Brown's team loses c
d	\$50 if Charlie Brown's team wins d	\$-50 if Charlie Brown's team loses d
e	\$50 if Charlie Brown's team wins e	\$-50 if Charlie Brown's team loses e

¹⁷ In the case where no objects are baseball games (and, so, none are learned to be baseball games), the bets won't, together, guarantee a loss. Rather, Linus will break even. To avoid this potential escape hatch, we could include in our Dutch Book a third bet based on Linus's credence that some object is a baseball game. As long as that credence is positive, he will lose in the scenarios that we are worried about—where no objects are baseball games. But we make the stakes quite small, so that it can be overwhelmed by losses from the bets in the text in cases where he wins this bet. Since this side bet is not an interesting part of the inconsistency, I'll confine discussion of it to this note.

f

\$50 if Charlie Brown's team wins f

-\$50 if Charlie Brown's team loses f

Linus will choose three rows from the second chart to go through with. For each row he chooses, he will then be betting on the object of that row at different prices. For instance, suppose Linus chooses to go through with the bets from the second chart on b, c, and f. Then he is committed to betting on b, c, and f at different odds, as can be seen by comparing the b, c, and f rows on the two tables. Call an object that he is committed to betting on at different odds a *focal object* (in our example there are three focal objects).

Our strategy, in order to guarantee a loss, is to choose one of the focal objects. For all objects other than that focal object, we give Linus the reverse bet. Continuing with our example above, but without loss of generality, suppose we choose c as the designated focal object. Then, we put Linus in the ledger for \$1 for each baseball game other than c. For each such game, we pay him \$100 if it is won by Charlie Brown's team and we pay him \$0 if it is not. This bet neutralizes all of our first quantificational bet with Linus except for the part that constitutes a bet on c. We then have him perform a similar neutralizing bet with regard to his second quantificational bet. For each of the three objects followed through on in the second quantificational bet other than c (that is, b and f), we mark in the ledger that we owe \$50 to Linus. He then pays us \$100 for each game that is won by Charlie Brown's team and \$0 for each that is not.¹⁸ This bet neutralizes all of Linus's second quantificational bet except for the part that constitutes a bet on c.

These quantificational bets neutralize all parts of the original bets except for the parts that constitute bets on the designated focal object, c. Linus is committed to thinking that game c is 1% likely to be won by Charlie Brown's team and 50% likely to be won by Charlie Brown's team. So, if it is won by Charlie Brown's team, Linus loses \$99 on the first bet and wins \$50 on the second for a net loss of \$49. If c is not won by Charlie Brown's team, Linus wins \$1 on the first bet and loses \$50 on the second, for a net loss of \$49. So, we've devised a system of bets that Linus's quantificational credences sanction as fair that guarantee him a loss.

Let's reflect on why we had to use the more complicated strategy here. Unlike in our original example with Charlie and the football, here we didn't know ahead of time precisely where the inconsistency would manifest itself. But we did know that it would manifest itself somewhere. So, we had a strategy that, no matter how Linus choose to bet based on his less-than-universal quantificational credence, guaranteed him a loss. In our particular case we supposed that the parts of the second quantificational bet that Linus followed through on were those on b, c, and f, and we chose c as the designated focal object for forming our Dutch Book. Of course, Linus could have chosen differently, but it should be clear that the same general strategy would work no matter which three objects he chose. Thus, he was guaranteed a loss even though we couldn't tell ahead of time where exactly that loss would come about (i.e., what the focal objects would be).

5.2. A Preliminary Note about Our Dutch Book Theorem for Quantificational Consistency

Before proving our theorem, we need to discuss the nature of the examples we have just

¹⁸ We might be worried that his quantificational credence need not sanction this bet as fair, though it did sanction as fair the reverse universal quantificational bet. This is because it is, in a sense, up to him which are the three baseball games to bet on. However, in making that choice with his first less-than-universal quantificational bet, he has already committed himself. Considering only agents who actually follow through on this commitment is, admittedly, an idealization. But this idealization is of a kind with some of those mentioned in Section 3, like the idealization that the agent not bet out of line with her credences when she sees a loss coming.

seen and of the theorem we will prove. In Section 3 we explored the standard characterization of a Dutch Book as a series of bets that an agent's credences sanction as fair which guarantee her a loss no matter which way the world turns out to be. For our purposes, we need to slightly amend this notion. For instance, suppose we are dealing with a quantificational bet sanctioned as fair by a quantificational credence of strength n , for q of the F s, that they are G s. For our Dutch Book Theorem, we are only going to be concerned with a subset of the ways that the world could turn out to be—those in which the only thing that an agent learns are what objects are all and only the F s.¹⁹ Admittedly, this restricts us to only a proper subset of the ways the world could turn out to be. In this subsection, I'll first explain why this restriction is necessary for our Dutch Book Theorem, and I'll then explain why the restriction doesn't diminish the support that the Dutch Book gives for Quantificational Consistency.

First, suppose that we have an agent who, at t_1 , has the above quantificational credence as part of a credal state that violates Quantificational Consistency. There are cases where she is not subject to a guaranteed loss because she learns some things at t_2 (before learning which objects are all and only the F s) that cause her to revise her quantificational credence, dissolving the irrationality. In this case, the quantificational bet sanctioned as fair at t_1 would be cancelled with no payout, even though our agent was indeed irrational at t_1 . The problem is that we don't want her uncertainty, at t_1 , about what objects are the F s to affect whether we find a way to show that she is irrational. By focusing only on ways the world could be where she learns only which objects are all and only the F s after the placement of the bet, we avoid situations like this.

But can we make such a restriction and still claim, as I will below, that a Dutch Book Theorem of this sort supports Quantificational Consistency?²⁰ First, we should keep in mind that Quantificational Consistency is a synchronic norm. So, our target for investigation is not the progression of a credal state over time. More importantly, the only reason time plays a role in our examples is that it allows uncertainty about which objects are all and only the F s to be resolved. If not for this, the Dutch Book could take place at a single time. Since the only reason we care about running time forward is for this purpose, restricting the ways the world could turn out to be to just those where what is learned is exactly which objects are all and only the F s does not take away from the evidential impact of these Dutch Books. We are not imposing restrictions on things that affect the payout in a suspect way (e.g., by ignoring possible worlds where some objects are F or some other objects are G). So, this move does not take away from the degree to which our Dutch Book Theorem supports our norm, Quantificational Consistency.

We could even think of the temporal factor of our Dutch Book as merely a dramatization of the formal or logical issue of the possible ways of extending a credal state from uncertainty (about which objects are all and only the F s) to belief. And the Dutch Book Theorem that I will prove in Section 5.3 shows that any of these possibilities being actualized leads to a loss for the agent. It is not meant to tell us anything important about how a credal state should or should not evolve over time (like a conditionalization norm does) nor about whether a current credal state is consistent with a future credal state (like a reflection norm does). Appreciating this allows us to see that a worry about the epistemic significance of Diachronic Dutch Book Arguments (like those in favor of conditionalization and reflection) does not apply here. For instance, Christensen (1991) gives a number of examples which aim to show that Diachronic Dutch Book Arguments, by themselves, fail to provide support for norms of epistemic rationality. He thinks that they fail

19 If a reader quibbles with my continuing to use the name "Dutch Book", she is welcome to think in terms of a nearby notion (what we could call a "Dutch Book*") which, as I'll argue, has the same evidential force.

20 Or, more carefully, that it does so on the condition that any Dutch Book Theorems offer such support.

because we don't have any antecedent reason to think that diachronic inconsistency is irrational (though he stresses that we do have reason to think that synchronic inconsistency is).

A major reason for doubting that Diachronic Dutch Books provide evidential support for norms on rationality is that one can be proven for a proposed principle of rationality that is generally thought to be false—the principle of reflection. I'm inclined to agree with those who, like Briggs (2009) and Mahtani (2015), maintain that the problem with the Dutch Book Argument for the principle of reflection is not fundamentally about either the principle or the series of bets being diachronic (see also Vineberg, 1997, for worries about the examples Christensen uses to come to the conclusion that we just saw). But even if Christensen is right, his worry does not apply to our (to be proved) Dutch Book Theorem for Quantificational Consistency, given the characterization that I gave in the previous paragraph.²¹ This is because there is reason to think that synchronic consistency is required between one's quantificational credences (and one's ordinary credences). Quantificational Consistency is not a norm on consistency of one's credal state across time or at different times. It is a norm prescribing synchronic consistency.²² We have, then, by Christensen's own lights, independent reason for thinking that the type of inconsistency that our Dutch Book will expose is, indeed, irrational—precisely what Christensen thinks makes Synchronic Dutch Book Arguments plausible. Thus, since our Dutch Book Theorem is not fundamentally diachronic—appearances of diachronicity are merely a dramatization of the resolution of uncertainty—the proof of the Theorem has exactly as much evidential import as the standard Dutch Books that we saw in Section 3.²³

5.3. Proof of the Dutch Book Theorem for Quantificational Consistency

We saw how to guarantee a loss for our protagonists in the examples above. Now we will see that in all cases where an agent violates Quantificational Consistency, she will be subject to a Dutch Book (i.e., the Dutch Book Theorem for Quantificational Consistency is true). The proofs of many Dutch Book Theorems give a recipe for forming bets that would guarantee a loss to the agent who violated the norm. The proof that I will give of the Dutch Book Theorem for Quantificational Consistency won't do that. Instead, I'll give a simple existence proof, showing that a Dutch Book exists in all cases where Quantificational Consistency is violated. Given the complexity of our quantificational bets and the different methods of Dutch Books corresponding to different ways of violating Quantificational Consistency, a mere existence proof is more streamlined. If the reader wants a recipe for forming Dutch Books in particular cases, she can glean one from the examples in the previous subsection.

Our first preliminary is to model the connection between credences—both ordinary and

21 Our Dutch Book Theorem also retains its force on Briggs's (2009) way of drawing the line between those Dutch Books that signal irrationality and those that don't. The latter, which include the Dutch Book Theorem for the principle of reflection, expose self-doubt, but this is not what is exposed by the one that I will prove in Section 5.3. I suspect that we would also be able to give a Dutch Book Theorem for Quantificational Consistency of the sort suggested in Mahtani (2015), though the details would require a much lengthier discussion about the structure of quantificational credences and what parts of that structure are and are not the sorts of things that require an interpretation (in the logical sense which is at play in her characterization of Dutch Books). See Footnote 9.

22 A norm prescribing synchronic consistency is stronger than simply a synchronic norm. As Christensen (1991: 240) points out, the principle of reflection is a synchronic norm, since it merely prescribes that one has a particular credence at a particular time. But it is not a norm prescribing synchronic consistency, since the prescription it makes for that credence is that it should relate in the right way to one's credence at a future time. Quantificational Consistency is both a synchronic norm and a norm prescribing synchronic consistency.

23 Thanks to an anonymous reviewer whose comments prompted the addition of this subsection.

quantificational—and betting prices that are sanctioned as fair. Though the final machinery might look complex, I'll build the picture from simple roots in order to make clear why each change is introduced. We know that an ordinary credence sanctions as fair bets at particular prices on some proposition. For instance, a credence of strength .3 that it will rain today sanctions as fair both buying and selling a bet of \$3 to win \$10 if it rains today and win \$0 if it does not rain today. We want to represent all of the bets sanctioned as fair not just by a single credence but by a set of credences. The easiest way to do this is to let the set of credences itself, C , represent these bets. The betting price sanctioned as fair for each credence is determined in the way of the rain example just discussed. A credence of strength n in p sanctions as fair both buying and selling a bet of $\$n$ to win $\$1$ if p is true and $\$0$ if p is false (as well as scaled versions of this bet).

However, it will be beneficial (for reasons that will be discussed) for us to represent the bets sanctioned as fair by a set of credences, C , in a more complex way. We will represent the bets with a set of pairs consisting of a set of beliefs and C :

$$|C| = \{ \langle B, C \rangle : B \text{ is a set of beliefs} \}.$$

That is, we take C and successively pair it with each possible set of beliefs. Each pair represents the bets sanctioned as fair by the set of credences of that pair (in the way discussed above) given the set of beliefs it is paired with. Since a set of beliefs does not affect the bets sanctioned as fair by a set of credences, the belief portion of the pair plays no substantial role. No information is gained or lost in moving from the previous modeling choice to this slightly bulkier one.

The reason we model the bets sanctioned as fair by sets of ordinary credences with belief-credence pairs is that it allows us to have a unified way of representing the bets sanctioned as fair by sets of ordinary credences and by quantificational credences. In order to properly represent the bets sanctioned as fair by quantificational credences—quantificational bets—we need to allow for interaction between a set of beliefs and a set of credences. I'll explain why this is so for a universal quantificational credence of strength n , for each F , that it is G . But the point extends to less-than-universal ones as well. Our quantificational credence sanctions as fair a quantificational bet. As we saw in Section 4, this bet is not equivalent to any particular ordinary bet or set of ordinary bets. So, we cannot simply represent it by a set of credences. The quantificational bet can be open-ended, in the sense that the agent need not know which objects are all and only the F s.²⁴ Once we add to the agent's belief state beliefs about which objects are all and only the F s, she is committed to bets on those individual objects—bets that are equivalent to those sanctioned as fair by a set of credences. But it is only the addition of these beliefs that commits her to betting in line with those odds. What that means is that we can represent this way that the quantificational bet sanctioned as fair by her quantificational credence plays out only as a set of ordinary credences, given some set of beliefs. Since there are a number of ways that the bet could play out, I group these pairs together and use a set of pairs of sets of beliefs and sets of credences to represent the quantificational bet sanctioned as fair by a quantificational credence.

While this general explanation helps to justify our modeling choice, it is also helpful to see exactly how we model the quantificational bet sanctioned as fair by a quantificational credence. Let's first think about our universal quantificational credence of strength n , for each F , that it is G , and then generalize to other cases. As we just saw, modeling the quantificational bet sanctioned as fair by this credence involves an interplay between beliefs and credences. So, for

24 Even if the agent does know which objects are all and only the F s, in the case of a less-than-universal quantificational credence, the agent need not yet know which q parts of the quantificational bet she will follow through on (i.e., which of these objects to bet on at odds n that they are G).

each set of beliefs an agent could have about which objects are all and only the Fs, there is a set of credences that sanctions as fair the ordinary bets that an agent is committed to. Let's first investigate the pair for one way of deciding which objects are all and only the Fs. Let b_1 be the belief that o_1 is F, b_2 be the belief that o_2 is F, . . . , b_j be the belief that o_j is F. Let $B^* = \{b_1, b_2, \dots, b_j\}$ and suppose that B^* is the set of beliefs about which objects are all and only the Fs. There is, then, a set of credences, C^* , that corresponds to B^* . Where c_1 is a credence of strength n that o_1 is G, c_2 is a credence of strength n that o_2 is G, . . . , and c_j is a credence of strength n that o_j is G, $C^* = \{c_1, c_2, \dots, c_j\}$. Thus, $\langle B^*, C^* \rangle$ is a member of the set of pairs that represents the quantificational bet, in the sense that for an agent who takes that bet and comes to have B^* as the set of beliefs about which objects are all and only the Fs, the quantificational bet requires the agent to make bets that are equivalent to those sanctioned as fair by C^* . However, one could have different beliefs (other than B^*) about which objects are all and only the Fs. For any set of beliefs about which objects are all and only the Fs, the universal quantificational bet will require an agent to take bets that are equivalent to a set of ordinary bets sanctioned as fair by the set of credences of strength n about each object believed to be F being G (in the way we just saw with B^* and C^*). We can model the universal quantificational bet by all possible pairs of this sort: $\{\langle B', C' \rangle: B' \text{ is a set of beliefs about which objects are all and only the Fs, and for each object, } x, \text{ believed in } B' \text{ to be F, there is a credence in } C' \text{ of strength } n \text{ that } x \text{ is G}\}$.

Things are only slightly more complex in the case of a less-than-universal quantificational credence of strength n , for q of the Fs, that they are Gs. In the universal case, we had, for an agent with the set of beliefs about which objects are all and only the Fs, $\{b_1, b_2, \dots, b_j\}$, that the quantificational bet prescribed bets equivalent to the ordinary bets sanctioned as fair by $\{c_1, c_2, \dots, c_j\}$. This is not so for the less-than-universal case, since our agent need only follow through on bets including q of o_1, o_2, \dots, o_j . What we know, then, is that letting $\{b_1, b_2, \dots, b_j\}$ be the beliefs about which objects are all and only the objects that are F, the quantificational bet will prescribe bets equivalent to ordinary bets sanctioned as fair by some combination of q of the members of $\{c_1, c_2, \dots, c_j\}$. But, since it is up to the agent to choose, we can't write in stone which combination this will be.²⁵ For instance, if q is exactly three, we know that the quantificational bet will prescribe bets equivalent to the ordinary bets sanctioned by exactly three of the members of $\{c_1, c_2, \dots, c_j\}$. So, our set of pairs will include (where $j \geq 7$) $\langle \{b_1, b_2, \dots, b_j\}, \{c_1, c_2, c_3\} \rangle$, $\langle \{b_1, b_2, \dots, b_j\}, \{c_1, c_4, c_6\} \rangle$, and $\langle \{b_1, b_2, \dots, b_j\}, \{c_4, c_5, c_7\} \rangle$, among others. That is, it will include multiple pairs that include the same set of beliefs—in contrast to our modeling of universal quantificational credences. More generally, we can model the quantificational bet sanctioned as fair by Q (where Q is a quantificational credence with any strength of quantification, q) as follows:

$|Q| = \{\langle B', C' \rangle: B' \text{ is a set of beliefs about which objects are all and only the Fs, and for } q \text{ of the objects, } x, \text{ believed in } B' \text{ to be F, there is a credence in } C' \text{ of strength } n \text{ that } x \text{ is G}\}$.²⁶

25 It's worth noting that our lack of knowledge of an agent's choice in an actual scenario won't be a problem for us, since the Dutch Book will exist in exactly the cases where no matter which choices the agent makes, she is subject to a sure loss.

26 Strictly speaking, I've introduced two $|\bullet|$ functions. The first function takes a set of credences as its argument and yields, as its value, a set of pairs—the first member of which is a set of beliefs and the second member of which is that set of credences. The second function takes a quantificational credence as its argument and yields the set of pairs of sets of beliefs and sets of credences in the way that we just saw in the text. But there is a reason that I use the same notation to talk about two different functions. It is that each function is supposed to model

Now that we have a way of modeling ordinary and quantificational bets, we can take care of one final preliminary before proving our Dutch Book Theorem. We'll be concerned not simply with whether the quantificational bets sanctioned as fair by an agent's quantificational credences can, on some way of coming to have beliefs about which objects are all and only the Fs, avoid a sure loss, but whether the quantificational bets sanctioned as fair by her quantificational credences can, on a way of coming to have beliefs about which objects are all and only the Fs that takes into account the beliefs she already has, avoid a sure loss. So, we're only going to want to investigate ways of filling out the quantificational bet that incorporate the beliefs that she already has. We can define, for the agent's (possibly incomplete) set of beliefs about objects being F, B^\wedge , the following:

$|Q|_{B^\wedge} = \{ \langle B', C' \rangle : B' \text{ is a set of beliefs about which objects are all and only the Fs, } B^\wedge \subseteq B', \text{ and, for } q \text{ of the objects, } x, \text{ believed in } B' \text{ to be } F, \text{ there is a credence in } C' \text{ of strength } n \text{ that } x \text{ is } G \}$.

Now we're in a position to prove our theorem—that if an agent's credal state violates Quantificational Consistency, then there exists a Dutch Book against her. We'll do this in three steps. Suppose that an agent A has a (possibly empty) set beliefs B^\wedge about whether objects are F, quantificational credences, Q_1, Q_2, \dots, Q_k , and a set of ordinary credences, C. The Dutch Book Theorem for Quantificational Consistency follows obviously from three Lemmas:

Lemma 1: If A's credal state violates Quantificational Consistency, then $|Q_1|_{B^\wedge} \cap |Q_2|_{B^\wedge} \cap \dots \cap |Q_k|_{B^\wedge} \cap |C| = \emptyset$.

Lemma 2: If $|Q_1|_{B^\wedge} \cap |Q_2|_{B^\wedge} \cap \dots \cap |Q_k|_{B^\wedge} \cap |C| = \emptyset$, then A is committed to betting on some proposition at different odds.

Lemma 3: If A is committed to betting on some proposition at different odds, then there exists a Dutch Book against A.

I will prove the Lemmas in reverse order.

Lemma 3 says that if A is committed to betting on some proposition at different odds, then there exists a Dutch Book against A. We can see this as follows (Skyrms 1986: 187): Imagine an agent for whom a bet on p is sanctioned both at odds n and at odds $m \neq n$. Suppose for now that $m < n$. Then we have the agent bet \$n on p to win \$1 if p is true and \$0 if p is false. And we buy from the agent a bet of \$m on p to win \$1 if p is true and \$0 if p is false. Such an agent is guaranteed a loss. Whether p is true or false, she wins $\$(m-n)$, and since $m < n$, this constitutes a loss. In the case where $m > n$, we simply switch places with the agent and a loss is guaranteed in the same way. This completes the proof of Lemma 3—that if A is committed to betting on some proposition at different odds, then there exists a Dutch Book against A.

Now let's suppose that $|Q_1|_{B^\wedge} \cap |Q_2|_{B^\wedge} \cap \dots \cap |Q_k|_{B^\wedge} \cap |C| = \emptyset$ toward showing that A is then committed to betting on some proposition at different odds (as Lemma 2 states). Remember that $|Q|_{B^\wedge}$ is a set of pairs. Each pair represents a set of ordinary bets that she would be committed to on some possible way of resolving her uncertainty about which objects are all and only the Fs (and making the choices left to her, in the case of a less-than-universal quantificational credence). So, the set $|Q|_{B^\wedge}$ as a whole represents, for each of the possible ways of resolving her uncertainty about which objects are all and only Fs (and making the choices left to her, in the case of a less-than-universal quantificational credence), the set of ordinary bets that she would be committed to. Now let's consider two quantificational credences, Q_1 and Q_2 (and let's adopt the convention that Q_y is a quantificational credence of strength n, for q of the F_y s that

the bet (or bets) sanctioned as fair by its argument. So, I've introduced this technical sloppiness for the purpose of conceptual unity.

they are $G_{y,s}$). $|Q_1|_{B^\wedge} \cap |Q_2|_{B^\wedge}$ represents, for each of the possible ways of resolving her uncertainty about which objects are all and only the F_1 s and which are all and only the F_2 s, the set of ordinary bets that would satisfy the betting commitments of both Q_1 and Q_2 . What is important for us to consider is the significance of cases where $|Q_1|_{B^\wedge} \cap |Q_2|_{B^\wedge} = \emptyset$. If this is the case, then, for each of the possible ways of resolving her uncertainty about which objects are all and only the F_1 s and which are all and only the F_2 s, there is no single set of ordinary bets that would satisfy the betting commitments of both Q_1 and Q_2 . Given B^\wedge , Q_1 commits her to bet at one of various odds on whatever objects she comes to believe to be all and only the F_1 s, that they are G_{1s} , and Q_2 commits her to bet at various odds on whatever objects she comes to believe to be all and only the F_2 s, that they are G_{2s} . But there is no single set of odds that, given B^\wedge , allows her to satisfy both betting commitments with respect to some proposition. So, given B^\wedge , Q_1 and Q_2 commit the agent to betting on that proposition at different odds. The same is true when we consider any finite number of quantificational credences. If $|Q_1|_{B^\wedge} \cap |Q_2|_{B^\wedge} \cap \dots \cap |Q_k|_{B^\wedge} = \emptyset$, then given B^\wedge , Q_1, Q_2, \dots, Q_k commit the agent to betting on some proposition at different odds. Finally, there are possible cases where $|Q_1|_{B^\wedge} \cap |Q_2|_{B^\wedge} \cap \dots \cap |Q_k|_{B^\wedge} \neq \emptyset$, but $|Q_1|_{B^\wedge} \cap |Q_2|_{B^\wedge} \cap \dots \cap |Q_k|_{B^\wedge} \cap |C| = \emptyset$. In such a case, given B^\wedge , Q_1, Q_2, \dots, Q_k do not commit the agent to betting on some proposition at different odds. But, given B^\wedge and C , they do commit the agent to betting on some proposition at different odds. That is, any way of, given B^\wedge , resolving one's uncertainty about the objects that are F_1, F_2, \dots, F_k , Q_1, Q_2, \dots, Q_k commit the agent to betting on some proposition at odds different from those sanctioned as fair by some member of C . This completes the proof of Lemma 2.

We can finish our proof of the theorem by proving Lemma 1—that if A 's credal state violates Quantificational Consistency, then $|Q_1|_{B^\wedge} \cap |Q_2|_{B^\wedge} \cap \dots \cap |Q_k|_{B^\wedge} \cap |C| = \emptyset$. Suppose that A 's credal state violates Quantificational Consistency. Then there is no complete extension of B^\wedge and complete extension of C that satisfy each of Q_1, Q_2, \dots, Q_k . We can represent this by saying:

$\{ \langle B^+, C^+ \rangle : B^+ \text{ is a complete extension of } B^\wedge, C^+ \text{ is a complete extension of } C, \text{ and } \langle B^+, C^+ \rangle \text{ satisfies each of } Q_1, Q_2, \dots, Q_k \} = \emptyset$.

Now what remains to be shown is that if this is true, then $|Q_1|_{B^\wedge} \cap |Q_2|_{B^\wedge} \cap \dots \cap |Q_k|_{B^\wedge} \cap |C| = \emptyset$. Given the truth of the antecedent, we know the following:

$\{ \langle B^+, C^* \rangle : B^+ \text{ is a complete extension of } B^\wedge, C^* \text{ is totally decided, and } \langle B^+, C^* \rangle \text{ satisfies each of } Q_1, Q_2, \dots, Q_k \} \cap \{ \langle B, C^+ \rangle : B \text{ is a set of beliefs and } C^+ \text{ is a complete extension of } C \} = \emptyset$.

This is because though the first set of pairs may include pairs the second member of which is a totally decided set of credences which is not a complete extension of C , the second set contains only pairs which have complete extensions of C as their second member. Now, the pairs of totally decided beliefs and totally decided credences that satisfy each of a set of quantificational credences are the pairs that satisfy the first quantificational credence and the second and the third, etc. So, we know the following:

$\{ \langle B^+, C^* \rangle : B^+ \text{ is a complete extension of } B^\wedge, C^* \text{ is totally decided, and } \langle B^+, C^* \rangle \text{ satisfies } Q_1 \} \cap \{ \langle B^+, C^* \rangle : B^+ \text{ is a complete extension of } B^\wedge, C^* \text{ is totally decided, and } \langle B^+, C^* \rangle \text{ satisfies } Q_2 \} \cap \dots \cap \{ \langle B^+, C^* \rangle : B^+ \text{ is a complete extension of } B^\wedge, C^* \text{ is totally decided, and } \langle B^+, C^* \rangle \text{ satisfies } Q_k \} \cap \{ \langle B, C^+ \rangle : B \text{ is a set of beliefs and } C^+ \text{ is a complete extension of } C \} = \emptyset$.

Remember our definition of satisfaction: Where q is any quantifier, we can say that a pair

consisting of a totally decided set of beliefs, B , and a totally decided set of credences, C , *satisfies* a quantificational credence of strength n , for q of the F s, that it is G if and only if for each of q of the objects believed to be F in B , x , C contains a credence of strength n that x is G . So, by explicitly substituting this in all but the final set, we have the following:

$\{ \langle B^+, C^* \rangle : B^+ \text{ is a complete extension of } B^\wedge, C^* \text{ is totally decided, and for each of } q \text{ of the objects believed to be } F_1 \text{ in } B^+, x, C^* \text{ contains a credence of strength } n \text{ that } x \text{ is } G_1 \} \cap$
 $\{ \langle B^+, C^* \rangle : B^+ \text{ is a complete extension of } B^\wedge, C^* \text{ is totally decided, and for each of } q \text{ of the objects believed to be } F_2 \text{ in } B^+, x, C^* \text{ contains a credence of strength } n \text{ that } x \text{ is } G_2 \} \cap$
 $\dots \cap \{ \langle B^+, C^* \rangle : B^+ \text{ is a complete extension of } B^\wedge, C^* \text{ is totally decided, and for each of } q \text{ of the objects believed to be } F_k \text{ in } B^+, x, C^* \text{ contains a credence of strength } n \text{ that } x \text{ is } G_k \} \cap$
 $\{ \langle B, C^+ \rangle : B \text{ is a set of beliefs and } C^+ \text{ is a complete extension of } C \} = \emptyset.$

Complete extensions of sets of beliefs and credences make them totally decided about all matters. But, the empty intersection of the sets just mentioned can only arise in virtue of beliefs and credences about objects which are, in the extensions, believed to be F_1, F_2, \dots, F_k . The sets place no constraints on beliefs and credences about other objects. So, an empty intersection results even if we simply look at extensions of B^\wedge that are not complete but merely include beliefs about whether all objects are F_1, F_2, \dots, F_k . Thus, we know the following is true:

$\{ \langle B', C' \rangle : B' \text{ is a set of beliefs about which objects are all and only the } F_{1s}, B^\wedge \subseteq B', \text{ and, for } q \text{ of the objects, } x, \text{ believed in } B' \text{ to be } F_1, \text{ there is a credence in } C' \text{ of strength } n \text{ that } x \text{ is } G_1 \} \cap$
 $\{ \langle B', C' \rangle : B' \text{ is a set of beliefs about which objects are all and only the } F_{2s}, B^\wedge \subseteq B', \text{ and, for } q \text{ of the objects, } x, \text{ believed in } B' \text{ to be } F_2, \text{ there is a credence in } C' \text{ of strength } n \text{ that } x \text{ is } G_2 \} \cap \dots \cap$
 $\{ \langle B', C' \rangle : B' \text{ is a set of beliefs about which objects are all and only the } F_{ks}, B^\wedge \subseteq B', \text{ and, for } q \text{ of the objects, } x, \text{ believed in } B' \text{ to be } F_k, \text{ there is a credence in } C' \text{ of strength } n \text{ that } x \text{ is } G_k \} \cap$
 $\{ \langle B, C^* \rangle : B \text{ is a set of beliefs and } C^* \text{ is a complete extension of } C \} = \emptyset.$

Given our modeling assumptions, this is equivalent to:

$|Q_1|_{B^\wedge} \cap |Q_2|_{B^\wedge} \cap \dots \cap |Q_k|_{B^\wedge} \cap |C| = \emptyset.$

Thus, we've shown that Lemma 1 is true—that if A 's credal state violates Quantificational Consistency, then $|Q_1|_{B^\wedge} \cap |Q_2|_{B^\wedge} \cap \dots \cap |Q_k|_{B^\wedge} \cap |C| = \emptyset.$

We've thus proved our three Lemmas. The Dutch Book Theorem for Quantificational Consistency follows from them by a couple of applications of hypothetical syllogism.

6. Conclusion

In this paper, we have seen that for any agent violating Quantificational Consistency, there is a Dutch Book. That is, there is a sequence of bets that can be made that will guarantee the agent a loss. Thus, I've shown the following conditional conclusion: if Dutch Book Theorems do support epistemic norms, then we have support for Quantificational Consistency.

In addition to explicitly establishing the above conclusion, I hope to have also established another task. Dutch Book Theorems have been proved for a number of purported norms involving ordinary and conditional credences. We've seen that they can also be proved for our relative newcomer in the theory of credal states—quantificational credences. The fact that a Dutch Book Theorem involving quantificational credences can be proved provides more evidence for accepting quantificational credences into our credal ontology and as a topic for study in formal epistemology. Furthermore, the discussion of quantificational bets (in addition to being interesting in its own right) has shown how quantificational credences are distinctive sorts

of credal attitudes. As we saw, they rationalize distinctive sorts of bets—and choices to act, in general—about which ordinary and conditional credences may be silent. In general, we've gotten a fuller picture of how quantificational credences fit into the landscape of credal attitudes.

Acknowledgements

Thanks to Mark Schroeder, Shyam Nair, Barry Lam, Justin Dallmann, and two anonymous reviewers for comments on and discussion of earlier drafts of this paper.

References

- Armendt, Brad (1992). Dutch Strategies for Diachronic Rules: When Believers See the Sure Loss Coming. *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, 217–229.
- Briggs, Rachael (2009). Distorted Reflection. *Philosophical Review*, 118(1), 59–85.
- Buchak, Lara (2013). *Risk and Rationality*. Oxford University Press.
- Buchak, Lara (2014). Risk and Tradeoffs. *Erkenntnis*, 79(S6), 1091–1117.
- Christensen, David (1991). Clever Bookies and Coherent Beliefs. *Philosophical Review*, 100(2), 229–247.
- Christensen, David (1996). Dutch-Book Arguments Depragmatized: Epistemic Consistency for Partial Believers. *Journal of Philosophy*, 93(9), 450–479.
- de Finetti, Bruno (1964). Foresight. Its Logical Laws, Its Subjective Sources (Henry Kyburg Jr. Trans.). In Henry Kyburg Jr. and Howard Smokler (Eds.), *Studies in Subjective Probability* (93–158). John Wiley & Sons, Inc.
- Hájek, Alan (2005). Scotching Dutch Books?. *Philosophical Perspectives*, 19(1), 139–151.
- Hájek, Alan (2008a). Arguments for – or against – Probabilism. *The British Journal for the Philosophy of Science*, 59(4), 793–819.
- Hájek, Alan (2008b). Dutch Book Arguments. In Paul Anand, Prasanta Pattanaik, and Clemens Puppe (Eds.), *The Oxford Handbook of Rational and Social Choice* (173–195). Oxford University Press.
- Hedden, Brian (2013). Incoherence without Exploitability. *Noûs*, 47(3), 482–495.
- Hellman, Geoffrey (1997). Bayes and Beyond. *Philosophy of Science*, 64(2), 191–221.
- Jackson, Frank and Robert Pargetter (1976). A Modified Dutch Book Argument. *Philosophical Studies*, 29(6), 403–407.
- Lewis, David (1999). Why Conditionalize? In *Papers in Metaphysics and Epistemology* (403–407). Cambridge University Press.
- Lennertz, Benjamin (2015). Quantificational Credences. *Philosophers' Imprint*, 15(9), 1–24.
- Lennertz, Benjamin (in press). Probabilistic Consistency Norms and Quantificational Credences. *Synthese*.
- Maher, Patrick (1992). Diachronic Rationality. *Philosophy of Science*, 59(1), 120–141.
- Maher, Patrick (1997). Depragmatized Dutch Book Arguments. *Philosophy of Science*, 64(2), 291–305.
- Mahtani, Anna (2012). Diachronic Dutch Book Arguments. *Philosophical Review*, 121(3), 443–450.
- Mahtani, Anna (2015). Dutch Books, Coherence, and Logical Consistency. *Noûs*, 49(3), 522–537.
- Ramsey, Frank (1964). Truth and Probability. In Henry Kyburg Jr. and Howard Smokler (Eds.), *Studies in Subjective Probability* (61–92). John Wiley & Sons, Inc.
- Skyrms, Brian (1986). *Choice and Chance: An Introduction to Inductive Logic* (3rd ed.). Wadsworth Publishing Company.
- Teller, Paul (1973). Conditionalization and Observation. *Synthese*, 26(2), 218–258.
- van Fraassen, Bas (1984). Belief and the Will. *The Journal of Philosophy*, 81(5), 235–256.

- Vineberg, Susan (1997). Dutch Books, Dutch Strategies and What They Show about Rationality. *Philosophical Studies*, 86(2), 185–201.
- Vineberg, Susan (2001). The Notion of Consistency for Partial Belief. *Philosophical Studies*, 102(3), 281–296.
- Vineberg, Susan (2016). Dutch Book Arguments. In Edward Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*. Retrieved from <http://plato.stanford.edu/archives/spr2016/entries/dutch-book/>