

Reconciling Meticulous Divine Providence with Objective Chance

Robert C. Koons

Abstract

Ontological randomness can be defined in terms of objective probability: an event is *random*, relative to a set of purposes or meanings, just in case its objective probability (in the circumstances) is other than *zero* or *one*. There is a prima facie tension between objective probability and divine providence: if God has providentially arranged for a particular event E to occur, then its objective probability would seem to be *one*, whatever worldly processes were apparently involved in its production. I will first show that this tension creates real problems for six models of how to combine worldly chance with divine providence: determinism, generalized Molinism, Thomism, divine luck, the multiverse, and van Inwagen's theory of chance in a world sustained by God. I will then develop two possible solutions to this problem, solutions that can be combined with any of the first three of the six models. The first solution was recently proposed by Alexander Pruss. Pruss relies on a theistic version of David Lewis's account of objective chance, in which chance corresponds to an ideal frequency intended by God. I raise an objection to Pruss's solution involving the gambler's fallacy and propose an alternative, which relies on three steps. First, the creation of some physical parameter (such as, for example, quantum probability amplitude), which can be identified with objective chance. Second, the introduction of a kind of 'divine command theory' of rational subjective probabilities or credences, requiring us to conform our subjective probabilities to our expectations of objective chance. Third, a divine intention to make actual frequencies conform approximately to these subjective expectations in almost all cases. In both my solution and Pruss's, a kind of divine overdetermination of events is involved: God contrives to bring about specific events via genuinely explanatory chancy mechanisms.

1. Introduction

In the philosophical literature there are essentially two ways of defining randomness (Eagle 2018): as a characteristic of a chancy process, and as a result with certain intrinsic characteristics

(algorithmic or Kolmogorov randomness). In this paper, I want to focus on the first way: an event is random just in case and insofar as it is the product of an objectively chancy process. By a chancy process, I mean one that has an objective probability of resulting in one of several alternative outcomes. This definition might be consistent with determinism, depending on our definition of objective chance: that is, it might be the case that a process is determined to have one specific result and yet also has an “objective chance” of having a different, counterfactual result. However, there is at least a *prima facie* tension between determinism and objective chance: it would seem reasonable to assign probability one to the result that is determined to occur and probability zero to all incompatible results.

Some quantum physicists and philosophers of physics hope to do without objective probability altogether. This includes Quantum Bayesianism or QBism (Caves, Fuchs, and Schack 2002, Fuchs 2010, Fuchs Mermin and Schack 2014). QBism builds on the work of earlier work of Jaynes, de Finetti, and others (Jaynes 1968, 2003, de Finetti 1972). The main problem for QBism lies in the interpretation of Born’s rule, which directs us to assign a certain probability to certain outcomes, given a known quantum wave function. Jaynes tried to rely exclusively on symmetry considerations to derive their probabilities from quantum theory. However, as Fuchs explains, quantum probabilities go beyond classical probability’s Principle of Indifference, since it constrains our judgments about both actual and counterfactual likelihoods (Fuchs 2010, 12)! In addition QBists face a dilemma. If probabilities are *merely subjective*—just autobiographical statements about our mental states—how can we “discover” probabilities by empirical study of external, physical processes? If, alternatively, QBists identify Born probabilities with the normative probability of an *ideal* agent, we confront the similar problem of explaining how we can discover a normative truth through empirical method (see Bacciagaluppi 2014). Fuchs even compares Born’s rule to the Ten Commandments (Fuchs 2010, 8-9)! My own proposal (to be laid out in section 5) can be thought of as a way of making sense of QBism: identifying Born probabilities with normative truths anchored in God’s intentions, and providing an account of how we can uncover facts about those divine intentions through empirical investigation.

Finding a satisfactory philosophical account of objective chance is a problem for everyone, but it is also a special problem for theists, especially theists who hold that God exercises a certain

degree of meticulous providence over creation, i.e., that God has in mind certain very specific, particular events that He intends, effectively, to bring about. Suppose, for example, that God intended for the astronomical, geological, and biological processes of creation to bring into being one particular human being, say Abraham, at a particular point in time. Since God is omnipotent, his intention could not fail to succeed. How, then, could Abraham's existence be, even in part, the product of chancy processes, processes with an objective chance of not resulting his existence (or the existence of any human beings, for that matter)?

There are two reasons for thinking this a serious question. First, it seems to be true that nature, as science reveals it to be, is filled with genuinely chancy processes. Quantum mechanics supports this idea in an especially acute fashion, since the Bell theorem rules out the most natural "ignorance" interpretation of quantum probabilities (that is, the existence of local hidden variables). Second, many branches of science, including statistical mechanics and evolutionary biology, rely on statistical explanations of observed phenomena, explanations that presuppose that the phenomena in question are the products of chancy processes. If human beings exist because God effectively intended that they, specifically and in particular, should exist, in what sense could statistical explanations in evolutionary biology also explain why such a species as humanity should exist?

In a recent unpublished paper (Pruss 2016), Alexander Pruss discusses five ways of reconciling meticulous chance and meticulous providence that fail—or, at least, that fail in the absence of significant supplementation. These five ways are: determinism, generalized Molinism, Thomism, divine luck, and the multiverse. We can also consider Peter van Inwagen's model for the existence of chance in a world sustained by God, which suffers from some of the same problems identified by Pruss. I will discuss these six failed reconciliations in section 2. Pruss's own solution is a theistic version of David Lewis's best-fit theory of probabilistic laws. I present Pruss's solution in section 3 and raise several objections to it in section 4. My own proposal appears in section 5: a divine command theory of rational credences, combined with the identifying of objective probability with a particular physical parameter (the square of the wave amplitude of the quantum wave function). I argue that this solution preserves the advantages of Pruss's account while avoiding my objections to it.

2. Six Failed Reconciliations

2.1 Determinism

We might first try a deterministic model of the universe. On this model, meticulous providence is easy to explain: God has simply to set the right initial conditions for the universe in order to obtain any possible history that he prefers. Given deterministic laws, his intentions are certain to succeed. But, as we saw, determinism seems *prima facie* inconsistent with objective chance. However, this inconsistency might be only apparent. As Pruss points out, classical (pre-quantum) thermodynamics made use of objective probabilities and statistical explanations, despite the fact that Newton-Maxwell dynamics were (almost) perfectly deterministic. Such classical statistical mechanics presupposes that we can identify objective probability with something like volume in a natural phase or state space: the larger the volume taken up by a set of states in that space, the greater its objective probability.

However, this underlying picture is inconsistent with meticulous providence. If God intentionally set the initial conditions of the universe in order to achieve a set of preferred outcomes, then there is no sense in which volumes of initial conditions that would lead to outcomes incompatible with God's intentions had any finite probability. Pruss asks us to imagine a perfectly skilled coin-flipper, who is able to produce Heads or Tails at will. If the flipper produces a sequence that is close to 50% Heads, then the only explanation of this fact must go through the flipper's actual intentions. The fact that the Heads-producing and Tails-producing sets of initial conditions are approximately equal in volume is completely irrelevant.

2.2 Generalized Molinism

Molinism is the theory (based on the work of Luis Molina) that God knows all of the "counterfactuals of freedom," despite the fact that human free choice is always the result of an indeterministic process. That is, if *C* fully describes the relevant features of a possible human free choice *F*, then God knows (from all eternity) whether or not it is true that, if *C* were to

obtain, F would result. Generalized Molinism extends such divine “middle knowledge” to the realm of chancy processes. As in the case of determinism, it is easy to use Molinism to explain meticulous divine providence: God can once again obtain any specific result He wants, so long as the result is *feasible* (i.e., actually obtainable via chance processes, given the actual truth-values of the relevant counterfactuals of chance) by simply fixing the right initial conditions. But, also once again, generalized Molinism fails to secure the reality of statistical explanation, for exactly the same reason that determinism fails to do so.

2.3 Thomism

We might reasonably suppose that the whole problem can be dissolved simply by relying on a central notion of Thomism: the distinction between *primary* and *secondary* causation. A result could be simultaneously chancy in the order of secondary causation (as produced by created causes) and completely determined in the order of primary causation (as specifically intended by God). I will argue in section 5 below that a variant of Thomism is part of the correct reconciliation, but Pruss points out an oddity that must be confronted.

For Thomists, the event that is C’s causing E (for any creaturely cause C and effect E) coincides by metaphysical necessity with the event of God’s willing that C cause E: any world containing one must also contain the other. Hence, if the objective chance of C’s causing E is x, then the objective chance of God’s willing that C cause E must also be x. Thus, we seem to be forced to attribute a kind of probabilistic propensity to God’s own volitions, as though God contained a kind of chancy causal mechanism, like an internal dice-throwing process, which is surely inconsistent with God’s aseity, simplicity, freedom, and perfection. It is surely the case that God acts indeterministically, but to project a mathematical measure onto God’s alternatives would seem to subordinate his decision-making process to something both internally complex and distinct from God’s essence. It is also implausible, as Pruss observes, that any such internal divine propensities would coincide perfectly with physically-based propensities discoverable by empirical science.

2.4 Divine Luck

On this model, God intends to bring about a particular event E . He sets up initial conditions that lead to a chancy process P , a process which has some probability of producing E spontaneously, and some objective probability of not doing so. God intends to intervene miraculously if P does not produce E spontaneously. If God is lucky, E will result from P , in which case E 's occurrence will have been overdetermined. If God's intentions are highly specific, and if the processes involved have propensities that are associated with probabilities significantly less than one, then God would have to be very lucky for this reconciliation to be successful.

2.5 Multiverse

The last model could be improved by adding many universes. With each additional universe, the chances of God's being sufficiently lucky in at least one of them improves. With enough universes, the chances of sufficient luck in at least one approaches certainty. This would work, but it makes it very unlikely that we inhabit a universe in which God's intentions are realized.¹ In addition, we might well suppose that God intends particular events to occur in each universe, in which case the existence of additional universes is irrelevant.

2.6 Peter van Inwagen's Model

Peter van Inwagen (van Inwagen 1988) argues that God can decree that the created world contains chancy processes, while simultaneously decreeing that these processes will eventuate in very specific outcomes. To simplify, suppose that there is just one process P , which undergoes a series of chancy transitions, T_1, T_2, \dots, T_n , with each T_i having a range of possible outcomes $E_{i,1}, E_{i,2}, \dots, E_{i,m}$ associated with objective probabilities $P(E_{i,1}), P(E_{i,2}), \dots, P(E_{i,m})$. These transition probabilities are particular, single-case facts about the outcomes—as we shall shortly see, they

¹ As Enis Doko has pointed out (in correspondence) this depends on assuming that each of the universes is equally real. God could create an infinity of simulated universes and then decide which of them to realize. This suggestion raises a new problem: what is it for a fully detailed simulation to become "real"? What are the unreal universes lacking? Can we tell that we have this indefinable element of "reality"?

are not fully determined by the underlying physical or psychological symmetries. Ordinarily, we would think that the probability of the occurrence of some final (n -stage) outcome $E_{n,j}$ would be the product of the probabilities, $P(E_{1,j}) \cdot P(E_{2,j}) \cdot \dots \cdot P(E_{n,j})$.

Pruss imagines that we can talk meaningfully about the internal structure of God's intentions. God intends that certain facts should obtain for the sake of certain other facts. In the case at hand, God intends certain particular facts in the mosaic for the purpose of making a certain system of laws (S) the best system of laws for the resulting world. God has intentions about what laws the world exhibits, and not just about individual events, taken one at a time.

Suppose that God intends a *disjunction* of final events, $(E_{n,1} \vee E_{n,2} \vee E_{n,3} \vee \dots \vee E_{n,m})$. Ordinarily, we would take the probability of this disjunction to be the sum of the probabilities $P(E_{n,1}) + P(E_{n,2}) + \dots + P(E_{n,m}) = \sum P(E_{n,i})$, which will be much less than one. However, van Inwagen imagines that God's decree can provide this disjunctive event with a probability of *one*. In other words, God can decree that some event in the set $E_{n,i}$ occurs, without decreeing which member of the set it is that occurs. God can leave it up to chance, in effect, leave it up to the chancy process P , to determine *which* member of the disjunction is actualized.

Van Inwagen does succeed in giving us a world in which there is both objective chance and a limited degree of meticulous providence. It is essential to van Inwagen's model that God does not decree every detail of history. He can decree that specific *types* of events (although not, perhaps, particular events) occur at particular junctures in the history of the world, while leaving it to chance *how* these event-types are brought about. There is, however, a serious drawback to van Inwagen's model. The actual objective probabilities depend in a very sensitive way to God's specific intentions and will therefore deviate in some (perhaps many) cases from the objective probabilities as we would ordinarily determine them in empirical science, i.e., from observed frequencies of similar setups. It is hard to see how empirical science can incorporate into the boundary conditions facts about divine intentions relating to the remote future. In addition, the van-Inwagen-objective-chance of a particular event would not always be determined solely by the volume of a corresponding region in a natural state space but would rather also depend on

which further events that event is likely to lead to and whether those further events are subject to God's decrees. This would seem to lead to a pervasive skepticism about objective chance.

Finally, we might reasonably suppose that God's decrees include the occurrence of particular events with particular participants and not just disjunctions of such particular events. For example, it seems plausible to suppose that God intended Abraham himself to exist, and not just Abraham or some Abraham-like counterpart. Such particular intentions would be incompatible with van Inwagen's model (except for intentions about the initial state of the universe).

3. Pruss's Solution: A Theistic Version of Lewis's Best-Fit Model

3.1 Lewis's Best-Fit Model

Pruss's new solution to the reconciliation problem builds on David Lewis's best-fit model of objective chance (Lewis 1980, 1994). Lewis's model was an extension of his own earlier work (Lewis 1973) on the Mill-Ramsey best-system theory of the laws of nature (Mill 1947, Ramsey 1978). According to the Mill-Ramsey-Lewis account of laws, a *law* is a theorem of the best axiomatic system of the particular natural facts of the actual world—the “Humean mosaic” of intrinsic qualities distributed across space and time. A system is *best* just in case it achieves the best combination of three values in relation to the actual mosaic: accuracy, comprehensiveness (strength), and simplicity.

The best-fit model of objective chance extends this model to include probabilistic laws. A probabilistic statement is a statement of objective chance (relative to the Humean mosaic of the actual world) just in case it achieves the best combination of intrinsic simplicity and fit to actual frequencies. The *degree of fit* between a probabilistic law and a corresponding frequency is simply a measure of the deviation between the two: the smaller the deviation, the closer the fit.

Lewis's best-fit model is a modification of the theory of frequentism of Hans Reichenbach and Richard von Mises (Reichenbach 1949, von Mises 1957). Frequentism identifies objective chance with long-run relative frequencies. The fundamental problem with the simple frequentist

model is that we expect there to be some deviation between objective chance and long-run relative frequency, especially if the relevant class is relatively small. We would not be surprised if it turned out that 50.0000001% of radium 223 atoms decayed in 11.43 days, even if the objective probability of decay in 11.43 days was exactly 50%. However, the frequentist must insist that objective probabilities *always* coincide *exactly* with relative frequency. On Lewis's best-fit model, this conclusion is not forced on us. We can trade a slight deviation for a simpler probabilistic law.

3.2 The Explanatory Weakness of Lewis's Chance

However, Lewis's best-fit model does inherit another central problem for frequentism: Lewisian objective chance cannot *explain* actual frequencies, since it ultimately *depends* on them. Suppose we observe a relative frequency F that is very close to the Lewisian best-fit probability r . Can we use the Lewisian probability to explain why F is close to r ? No, because the fact that F is close to r is part of the metaphysical explanation of why there is a probabilistic law assigning r (and not some other number) to the relevant class of events. To use Lewisian probabilities to explain statistical frequencies would thus be viciously circular.²

Here is where theism can help, as Pruss observes. Let's say that we have a probabilistic law of nature assigning an objective chance r to some class of outcomes E just in case God intends for the frequency of E to be close to r , as close as possible given his other aims and constraints. In other words, let's suppose that God intends for S (a system of laws, both deterministic and probabilistic) to be the best system of laws for the world as it actually comes to be. If it is a theorem of S that event E has probability r , then r is in fact E 's objective chance of occurring.

In Pruss's revised model, we can use objective chance to explain actual frequencies. The frequencies are (typically) close to value of the corresponding objective chance, and they are

² Here Pruss and I are rejecting accounts (like those of Loewer 2012) that draw a sharp separation between scientific and metaphysical explanation. The two are probably distinct, but it is hard to accept mixed cases of circularity, in which the fact that p scientifically explains the fact that q , while the fact that q metaphysically explains the fact that p . Realists about explanation have to suppose that any case of an explanatory relation involves a form of real, asymmetric dependency. (Thanks to Aaron Segal for bringing this to my attention.)

close to those value *because* the values represent objective chance, since God arranges things so as to make the fit as close as possible. The value of the objective chance depends on God's intention, not on the actual frequencies. The actual frequencies, in turn, depend on the chances.

3.3 Saving the Principal Principle

Pruss's revision also solves a serious problem that Lewis noted (Lewis 1980) with his own best-fit model: it comes into conflict with a widely accepted principle that constrains the relationship between rational credences and objective chance, the "Principal Principle" of probability.

The Principal Principle. $\text{Credence}(E/[H \ \& \ \text{Chance}(E) = r]) = r$

Let's suppose that r is significantly greater than 0, for some event-type E . Let E^* represent a very large and improbable ensemble of n occasions for E -type events in the future (beyond the scope of H), in which the relative frequency of E -type events is much lower than r —for simplicity's sake, let's set it at *zero*. The chance of E^* 's occurring should be small but finite, something like $(1 - r)^n$, assuming independence. Now, apply the Principal Principle. We can infer that our credence in E^* , conditional on $\text{Chance}(E^*) = (1 - r)^n$ and H , should itself be $(1 - r)^n$.

However, given the best-fit theory, it seems that E^* is actually inconsistent with $\text{Chance}(E^*) = (1 - r)^n$. It is metaphysically impossible for both to be true. In a world in which E^* occurs, the relative frequency of E must be much lower than r , since the actual frequency of E in such world must be far less than r . The laws of probability ensure that the probability of one proposition conditional on a proposition inconsistent with it, must be *zero*. Hence, the credence of E^* , conditional on $\text{Chance}(E^*) = (1 - r)^n$ and H , must be zero. But $0 \neq (1 - r)^n$ and H . Contradiction.

Pruss's model differs from Lewis's in this respect. It is *not impossible* in Pruss's account for the frequency of E and the chance of E to be far apart. Pruss's model stipulates that God must intend

to make the frequency of E as close as possible to the chance of E, *given God's others aims and intentions*. It is certainly conceivable that in certain cases God might have overriding reasons, reasons that would lead him to permit a wide deviation of frequency from chance. We might even be able to conceive a world in which every frequency deviates widely from its objective chance.

Lewis thought that he had overcome this problem (or “bug”) by focusing on the “admissibility” of the proposition $\text{Chance}(E^*) = (1 - r)^n$ (Lewis 1992). The Principal Principle can be applied only if the information on which the credence of E^* is being conditioned is *admissible* at the time at which it is being applied. That is, we cannot condition on a proposition that contains (even implicitly) future information relevant to the occurrence of E^* . But, given the best-fit model of chance, that is just what $\text{Chance}(E^*) = (1 - r)^n$ does—it implicitly provides information about the *future* frequency of E, since a proposition encoding an objective chance is covertly a proposition about a global relative frequency (including the future).

But, as Lewis recognized, to his temporary dismay (Lewis 1992, 485-6), this seems to make any application of the Principal Principle fallacious, given the constraint on inadmissible information and the best-fit theory of chance. Lewis argued (Lewis 1992, 486-7) that he could get around this by seeing that admissibility is a matter of degree. The Principal Principle is never strictly and exactly correct, but it can be *approximately* correct, so long as the proposition about chance does not provide *too much* information about the future. And that is exactly what the proposition that $\text{Chance}(E^*) = (1 - r)^n$ does in our present case, explaining the total failure of the application of the Principal Principle.

This was an ingenious solution but ultimately an unsatisfying one. As Lewis admitted, the Principal Principle is central to our concept of objective chance. Such a constitutive principle must be exactly correct—mere approximation is just not enough. Lewis's approximate solution is a bug, not a feature.

3.4 Pruss's Reconciliation of Providence and Chance

Pruss's model can be fruitfully combined with three of the attempted reconciliations: determinism, generalized Molinism, and Thomism. I prefer the combination of Pruss's model with Thomism. As Pruss points out, his model resolves the oddity that we noted above: the fact that the objective chance of an event's occurrence corresponds with the objective chance of a corresponding divine intention. Now we can ask: what is the truthmaker for the claim that the objective chance of God's intending E on occasion C is r ? The answer is this: the divine intention has chance r because God intends that the frequency of such divine intentions be as close to r as is possible. This clearly does not involve attributing to God some peculiar, sub-personal machinery within his decision-making process. Hence, the oddity is resolved in a satisfactory manner.

The Pruss-Thomist model can now reconcile meticulous providence with objective chance quite easily. We can now see why it is possible to explain a particular event (like the existence of Abraham) *both* as the result of an effective divine intention *and* as the result of a certain chancy processes. God intended (and caused it to be the case, in a primary mode) that Abraham's existence be explained in terms of secondary causation, including statistical explanations involving objective chances. Objective chances do really explain actual results, via God's intentions that they should do so (i.e., his intentions that the actual frequencies should approximate chances as closely as possible).

4. Some Objections to Pruss's Account

Pruss's account is clearly an improvement over Lewis's, and I believe that it is at least on the right track. Nonetheless, there are two problems or apparent problems, which should motivate us to look for a revised model.

4.1 The Gambler's Fallacy

The Pruss model would seem to license a version of the Gambler's Fallacy. Suppose that I know that there are only k possible occasions for the occurrence of an event of type E, and suppose that I have observed the first $k - 1$ occasions. Suppose further that, on these first $k - 1$ occasions, an

E-type event has occurred exactly $k/2$ times. Thus, I know that the relative frequency will be very close to $\frac{1}{2}$. Given the value of simplicity, that gives me good reason to think that the objective chance is exactly $\frac{1}{2}$, that is, that God has intended for the relative frequency to be as close to $\frac{1}{2}$ as possible. If an E-type event occurs on the last occasion, the frequency will be somewhat over $\frac{1}{2}$ --it will be $\frac{1}{2} + 1/k$. If instead a non-E-type event occurs, the frequency will be exactly $\frac{1}{2}$. Thus, I have good reason to expect that we will *not* see an E-type event on the last occasion, even though the objective chance for the occurrence of such an event is $\frac{1}{2}$. This reason need not be conclusive—*any reason at all* to prefer the non-occurrence of the E-type event to its occurrence on the last occasion is sufficient to falsify Pruss's model.

The Pruss model might be salvaged if we could identify a higher-order law or regularity that applies in this case. The first thing to note is that we should distinguish between objective chance and objective probability. An objective probability is a chance only when it is the conditional probability of an event-type conditional on *all causally prior facts* (see Pearl 2000 for details, especially chapters 1 and 2). So, in the example above, we need to consider the objective probability of an E-type event occurring on the last occasion, given that it has already occurred $k/2$ times on the previous $k - 1$ occasions. This defines a new event-type, which we can call type E^+ . Given the hypothesis, there is only one possible occasion on which an E^+ -type event can occur, so its relative frequency must be either 0 or 1. However, we might be able to find a more general class of event-types, call it F, that subsumes E^+ along with a large number of other, relevantly similar event-types. The objective chance of the occurrence of an F-type event will also be $\frac{1}{2}$, so God will have good reason to make the relative frequency of F-type events as close to $\frac{1}{2}$ as possible. Once I realize that the E^+ -type event is a member of this F class, I have good reason to anticipate its occurrence with a credence of exactly $\frac{1}{2}$, as required to avoid the Gambler's Fallacy.

Nonetheless, there still seems to be some grounds for being biased against the occurrence of an E-type event on this last occasion, given the value of matching perfectly a simple probability. But any bias will lead to a rational deviation of subjective probabilities from known objective chance, in contradiction to the Principal Principle.

4.2 The Credence/Chance Conceptual Gap

Finally, we can ask whether the Lewis-Pruss model is able to explain the normative bite that the Principal Principle represents. Why is it rational for us to apportion our credences according to the weights of objective chance? For both Lewis and Pruss, objective chance corresponds (at least approximately) to long-run, global relative frequency. But why should my subjective probability about any particular event correspond to global, long-run relative frequencies of similar events in similar circumstances? As John Maynard Keynes is supposed to have quipped, “in the long run, we’re all dead.” What would be *irrational* about setting my subjective probabilities about particular cases in a way that disregards such long-term facts and symmetries?

5. A Divine Command Theory of Rational Credence

5.1 The Model and its Advantages

Robert M. Adams’s divine command metaethics built upon earlier work in philosophical semantics by Keith Donnellan, Saul Kripke, and Hilary Putnam (Donnellan 1966, Kripke 1972, Putnam 1975), work which demonstrated the existence of necessary truths that are neither analytic nor knowable a priori. For example, it is a necessary truth that Venus is identical to Venus, and so it must also be a necessary truth that the Morning Star is identical to the Evening Star, since both phrases are simply names of Venus (Kripke 1972, 97-105). Similarly, since water is necessarily identical to water, water must be necessarily identical to H₂O, since both ‘water’ and ‘H₂O’ are names for the very same substance (Putnam 1975, 196-290). Nonetheless, these truths are not analytic or knowable a priori. No amount of reflection on the meaning of ‘the Morning Star’ or our concept of water could ever have led to the discovery that the Morning Star is the Evening Star, or that water is H₂O. These discoveries were empirical, learned *a posteriori*. Thus, we have *a posteriori* necessities and identities.

In a similar way, Adams proposed that the property of being morally wrong is identical to the property of being forbidden by God.³ Adams does not suppose that we can infer this identity by mere reflection on our concept of moral wrongness. The identity is discovered through a kind of theological and metaphysical inquiry that could be labeled ‘a posteriori’ in relation to metaethics. Despite this conceptual novelty, Adams proposed that the property we are in fact thinking of when we think of moral wrongness is the property of being forbidden by God.

I propose adapting Adams’s metaethics to the case of a certain cognitive or intellectual deontology: i.e., the *rational necessity* of conforming our subjective credences to certain normative principles. In our intellectual lives, as in our moral lives, we encounter certain categorical imperatives (to use Kant’s phrase): things that we must do or not do, regardless of their consequences in particular circumstances. We ought always to avoid logical inconsistency, and we ought to modify our credences in order to bring them in conformity to standard axiomatizations of probability (such as Kolmogorov’s or Popper’s). And, to come to the present case, we ought to conform our credences to our rational expectations of objective chance. On my theory of meta-normativity, these rational imperatives are in fact divine commands—things we are commanded by God to do in our intellectual lives.

I am assuming, for this model, that the relevant credences are subject to our voluntary control—that they consist in our making certain judgments of probability. Once we see that our judgments of probability are in conflict with the axioms of probability or with the Principal Principle, we are *obliged* (in a special, non-moral sense) to alter them in order to avoid the conflict.

How are these commands promulgated by God and known by us? Not, of course, by being carved in stone on Mt. Sinai. Rather, they are promulgated by being incorporated into certain normal operations and inclinations of the human mind. In this way, atheists and agnostics can be aware of the normative facts, without correctly understanding their metaphysical basis. In this respect, the laws of correct probabilistic thinking are like the natural moral law of Thomas Aquinas (see *Summa Theologiae* I-II, q90, a4).

³ Adams actually writes ‘forbidden by a loving God’, but since I assume that God is necessarily loving, I can omit this qualification.

In order to connect rational credence with objective chance, we have to suppose that objective chance corresponds to some real (possibly physical) parameter. In other words, God's command is that we apportion our credences to correspond to this chosen parameter. Since God is rational and benevolent, he has good reason to make the relative frequencies match the objective chance as closely as possible, since otherwise he would be issuing general commands that would lead rational agents to act sub-optimally in the long run.

What is this special parameter? In classical mechanics, it would correspond to the volume of an event in a natural state space. In quantum mechanics, there is an even simpler and more concrete parameter: the square of an event's quantum wave amplitude.

Thus, the model has a three-step structure:

- (A) God creates a special physical parameter (e.g., wave amplitude).
- (B) God commands that all rational creatures apportion their credences in accordance with some fixed function of that parameter (e.g., the square of the amplitude—the amplitude times its complex conjugate).
- (C) God has good reason to make the corresponding relative frequencies fit the rational credences as closely as possible, so that rational creatures who conform to the divine command would act optimally in the long run.

As in Pruss's model, my model can use objective chance to explain actual frequencies, thanks to step C of the model. Step B clearly closes the chance/credence conceptual gap. The model also avoids the gambler's fallacy, since we have good reason to conform to sound probabilistic principles (in order to conform to divine commands), and God has good reason to make frequencies optimal for rational agents in all circumstances, including the peculiar ones outlined in my scenario. Finally, there is no problem with the Principal Principle, since the correspondence of credence and chance is guaranteed immediately by the identity of chance with divine commands.

Why is step (A) necessary? Couldn't God have simply issued commands concerning our rational credences, without introducing a particular physical parameter? (Thanks to Aaron Segal for raising this point.) In my model, step (A) is needed to provide the particular content of God's commands in step (B). Here's an analogy. Suppose God commanded us to love our neighbor, i.e., to aim at promoting our neighbor's welfare. Such a command presupposes that there is such a parameter as individual welfare. In a similar way, step (B) presupposes that there is some variable, physical parameter upon which our rational credences are supposed to be based.

5.2 Objections

First, one might object that any divine command theory of normativity suffers from a vicious circularity. We would have to assume that there is a norm enjoining us to obey God's commands, but how can such a norm exist if all norms depend on God's commands? Robert Adams considered this objection in his essay, and he responded that his theory does not need any deontic norm directing us to obey God: it is sufficient if we have good reason to value such obedience. Not all reasons to act are constituted by deontic norms: there are also non-normative values to consider. In the case of God's commands, there are many reasons, independent of both morality and cognitive normativity, for valuing obedience. We value a good relationship with God, and, given the asymmetry in knowledge and character, such a good relationship depends on our obedience to his commands. Given God's creation of us and his subsequent generosity, we value our obedience as an expression of gratitude. It is aesthetically fitting that we should defer to God's commands, given the ontological asymmetry involved.

None of these reasons for obeying God's commands need be active in cases in which we feel bound by cognitive norms. It is sufficient that there exist good reasons to conform to those norms, whether or not we grasp what those reasons are. It is enough if we grasp the somewhat inchoate fact that there must be some good reason for us to conform to the norms we recognize, like the Principal Principle.

Second, there are grounds for worrying that my step C will not apply to cases that are beyond all human knowledge and concern. God's benevolence for us may give him reason to make relative

frequencies stick close to objective chances within the bounds of human knowledge and concern, but what could motive him to do so beyond those bounds? In response, I could argue that human beings do form beliefs in the form of unbounded, global generalizations. Physicists may well form the belief that the cosmic relative frequency of physical events matches closely the probability amplitude of those events. If we assume that God cares about whether we believe or have high confidence in truth or falsehood, regardless of whether we are ever able to verify these beliefs empirically, and regardless of whether these beliefs are of any practical import to us, then God does have sufficient reason to bring all relative frequencies close to the corresponding objective chances.

Third, Jeff Koperski has raised (in correspondence) the following worry. What can I say about people who are ignorant about the relevant divine commands? Didn't people assign probabilities rationally (or irrationally) prior to the discovery of quantum mechanics, and even prior to the discovery of classical statistical mechanics? Certainly, they did. Remember, first, that I am building on Robert Adams's account of divine command theory, which is explicitly a theory of *a posteriori* identity. Probabilistic rationality and irrationality do not depend on being aware of God's epistemic commands as such (i.e., under that theological description). Moreover, one cannot be even *materially* (so to speak) in violation of God's commands relating to quantum wave amplitudes without being aware of those amplitudes. Thus, the discovery of quantum mechanics involved the uncovering of new norms, norms that are as a matter of metaphysical fact (but not as a matter of *a priori* intuition) grounded in divine intentions. Prior to the discovery of the physical foundation of statistical mechanics, people could still violate other norms of probability (such as those encoded in the Kolmogorov axioms), but obviously they could not act contrary to God's intentions vis-à-vis quantum wave amplitudes or state space volumes. Progress in normative knowledge is possible in empirical science, just as it is possible in moral or political theory. As I mentioned in the Introduction, my proposal can be seen as providing metaphysical foundations for the similar claims made by Quantum Bayesians.

Bibliography

- Adams, Robert M. (1979), "Divine Command Metaethics Modified Again," *Journal of Religious Ethics* 7(1):66-79.
- Donnellan, Keith (1966), "Reference and Definite Descriptions," *The Philosophical Review* 75:281-304.
- Eagle, Antony (2018), "Chance versus Randomness," *The Stanford Encyclopedia of Philosophy* (Spring 2019 Edition), Edward N. Zalta (ed.), forthcoming URL = <https://plato.stanford.edu/archives/spr2019/entries/chance-randomness/>.
- Kripke, Saul (1972), *Naming and Necessity*, Cambridge, MA: Harvard University Press.
- Lewis, David K. (1973), *Counterfactuals*, Cambridge: Harvard University Press.
- (1980), "A Subjectivist's Guide to Objective Chance," in *Philosophical Papers*, volume 2, Oxford: Oxford University Press, pp. 83–132.
- (1994), "Humean Supervenience Debugged," *Mind*, 103: 473–90.
- Mill, John Stuart (1947), *A System of Logic*, London: Longmans, Green and Co.
- Pearl, Judea (2000), *Causality: Models, Reasoning, and Inference*, Cambridge: Cambridge University Press.
- Putnam, Hilary (1975), *Mind, Language, and Reality: Philosophical Papers, Volume 2*, Cambridge: Cambridge University Press.
- Pruss, Alexander R. (2016), "Evolution and Creation: Best-Fit Probabilities and Omnirationality to the Rescue," unpublished paper, November 1, 2016.
- Ramsey, Frank (1978), *Foundations*, London: Routledge and Kegan Paul.
- Reichenbach, Hans (1949), *The Theory of Probability*, Berkeley: University of California Press.
- van Inwagen, Peter (1988), "The Place of Chance in a World Sustained by God," in *God, Knowledge, and Mystery*, Ithaca: Cornell University Press, pp. 42-65.
- von Mises Richard (1957), *Probability, Statistics and Truth*, revised English edition, New York: Macmillan.