

Introduction

Philosophical theorising about God often begins with the idea, made famous by Anselm, that God is the greatest possible being. On such a view, God must have all the ‘great-making’ properties, or perfections. Some of the properties often thought to be great-making include *omnipotence* (being all-powerful), *omniscience* (being all-knowing), *omnibenevolence* (being wholly good) and perfect freedom. This philosophical approach to God is known as *perfect being theology* and the present work is situated in that tradition. There are different ways of working this project out, but for the most part, the finer details do not concern us. That’s because in what follows we are concerned, not with the internal coherence of the divine attributes, but with whether certain aspects of God’s being and activity are compatible with human free will.

In contemporary philosophy free will is often defined as *the control required to be morally responsible*. Here ‘control’ is to be understood as *control over our decisions, actions and (some of) their consequences*. This definition is functionalist inasmuch as it identifies free will by the role it plays. Throughout this Element we combine this definition of free will with the substantive idea that free will consists in having a choice about something. One reason for starting with this substantive notion of free will is that it is intuitive: all of us, through making choices, direct our lives to some degree or other. From choosing what to have for breakfast to which hobby to take up – or from whom to vote for to which friendships to invest in – having choices about these things is prized, even when some such choices are tough to make.

Indeed, even when our freedom of *action* is curtailed – when, for example, we are forced to do or to undergo something we might not want to – freedom of *choice* often remains; when it does, such freedom can alter, and bestow value on, our experiences. Henri Nouwen captures this nicely in a book of spiritual reflections:

Joy is what makes life worth living, but for many joy seems hard to find . . . [Yet] strange as it may sound, *we can choose* joy. Two people can be part of the same event, but one . . . may choose to trust that what happened, painful as it may be, holds a promise. The other may choose despair and be destroyed by it. What makes us human is precisely this freedom of choice. (1996: 37, emphasis added)

Moreover, making choices is often valued, not just because of the control over our lives it gives us, but also because in making choices we both *reveal something about* and *help to form* our identity as persons. As Dumbledore once explained to Harry Potter, “It is our choices, Harry, that show what we truly are, far more than our abilities” (Rowling 2015: 352).

Gary Watson goes so far as to say that a human “who never engaged in such activity [i.e. deciding] would be an agent only in a truncated sense” (2004b: 126). Watson’s point here is that there is a significant form of human agency which depends on being able to make reasoned choices. Moreover, this form of agency plays a central role in our conception of ourselves. Tillmann Vierkant, Julian Kiverstein and Andy Clark note the practical importance of choice when they write that the idea that humans can make autonomous decisions is “absolutely central to many of our social institutions, from criminal responsibility to the markets, from democracies to marriage” (2013: 1).

Another reason for starting with this substantive, choice-based understanding of free will is that the most venerable theological puzzles concerning God and free will only arise – or at least, arise in their most difficult forms – given this understanding of free will. Thus, if the existence and nature of God can be shown to be compatible with *this* conception of free will, it is a safe bet that whatever the precise nature of free will turns out to be, it will be compatible with the existence and nature of God.

Free will, or freedom – in this work we use the terms interchangeably – is a significant topic in many areas of philosophy of religion and theology. Freedom is relevant to the doctrine of God first and foremost because it is usually seen as one of the great-making properties that God must possess. How we understand God’s own freedom affects how we understand His activity in the world, and so informs statements of the doctrine of creation. The topic of God’s freedom also impacts theological anthropology, especially in the Judeo-Christian tradition, according to which human beings are said to be made in the image of God. For many theorists, a key part of this idea is that humans were given some measure of free will because God Himself is free. Free will is also a central topic of concern in the Christian doctrine of salvation. That’s because, according to the dominant view within the Western Church, humankind was subject to the Fall, an event through which evil entered human life, God’s creation became disordered, and human beings became unable to restore themselves to their original state. On this view God must intervene to save humans and salvation is therefore seen as a free gift from God – an act of His free will. At the same time, however, Christian theology usually teaches that humans are responsible for their state of brokenness. These two facts combine to generate the problem of grace and free will. The Eastern tradition within Christianity does not have quite the same doctrine of the Fall, and so doesn’t generate the same tension between grace and free will, but nevertheless has an important place for human freedom in its theology.¹

¹ Most of those whose work we discuss write from within the Christian tradition, and some of their positions – such as those related to the issues of sin, grace and salvation – arise out of a commitment to Christian orthodoxy. This commitment is most evident in Section 2, since

This Element focuses on two topics regarding which human freedom has significant relevance for the theistic worldview. The first is the problem of divine foreknowledge. In brief the problem is as follows: if God knows in advance how we will make each choice we face in our lives, and God cannot be wrong (so that what He knows must come to pass), how can our choices be free? This problem, which has been debated for millennia, is fascinating in part because it touches upon many philosophically difficult topics: causation, time, the future, truth and knowledge – in addition, of course, to God and freedom. For this reason the problem still garners much attention today among both theist and atheist philosophers. This is the topic of Section 1.

In Section 2 we consider the relationship between divine providence and human freedom. We understand *divine providence* to be God's acting to realise His ends. We construe this broadly so as to include, to use traditional theological language, God's *preserving* in being everything He creates, His *concurrence* with created causes, His *general provision* for what He creates and His *special or particular actions* within history. Sometimes the former two concepts are treated as part of the doctrine of creation, and the latter two are grouped together as God's *governance* of creation. For convenience, we treat them all as aspects of divine providence. The worry that arises in connection with free will here is that if God has complete control over all He has created – if God is sovereign over absolutely everything that comes to pass (as would seem fitting) – then there appears to be little room for any human *agency* at all, let alone human *freedom*. The challenge, as many theist thinkers see it, is to work out a doctrine of providence which takes seriously human freedom (to secure human responsibility) while also doing justice to God's sovereignty and control over His creation.

1 Divine Foreknowledge and Human Freedom

1.1 The Problem of Divine Foreknowledge and Human Freedom

The problem of divine foreknowledge and human freedom can be stated informally as follows. God is omniscient, and so knows everything there is to know, including facts about the future. God is also infallible, and so cannot be mistaken in what He knows. This means that God has always known how everyone will decide and act at any future time. But if God has always known just how each person will decide in the future, then those future decisions appear to be necessary or fixed in some way. And if those decisions are necessary or

many of those writing on divine providence seek to do justice to various biblical texts and confessional creeds which speak to God's control over human affairs.

fixed, then they are not free, because someone makes a free decision only if she is able to decide otherwise.

This venerable philosophical problem has been discussed in the Western tradition for thousands of years and has generated a vast literature. Part of the reason for that, of course, is that it concerns a central aspect of the doctrine of God and is of interest to every generation of theists. But the philosophical puzzle is also intriguing, and attracts the attention of many non-theists, because it touches upon so many other topics of philosophy: truth, the future, time and temporal ontology, modality and, of course, free will.

The argument outlined in the first paragraph of this section is, at its core, an argument which aims to show that exhaustive divine foreknowledge and human free will are incompatible. We understand *the problem of divine foreknowledge and human free will* to be the problem of either (a) saying where the reasoning of that argument goes wrong, or (b) justifying the rejection of divine foreknowledge or human freedom. We call those who pursue the first strategy *foreknowledge compatibilists*: they think there is a flaw in the argument and that divine foreknowledge and human freedom are compatible.² Those who pursue the second strategy are *foreknowledge incompatibilists*: they defend the validity of the reasoning in the argument and so must reject either divine foreknowledge or free will.

Before formally stating the argument that is discussed in what follows, we lay out the argument's assumptions. The argument assumes that God exists and is essentially eternal, essentially omniscient and infallible. There are two ways of understanding God's eternity. The first is that God is temporal and exists at all times. The second is that God is atemporal in the sense that He is 'outside' of time and doesn't have any temporal properties. Sometimes writers reserve the term 'everlasting' to refer to the first notion and reserve the term 'eternal' to refer to the second. However, given that both 'everlasting' and 'eternal' are widely found in English translations of the Bible, we prefer to avoid associating these terms with technical concepts and instead simply to be explicit about the concepts in play. The argument presented in what follows assumes that God is eternal in the first sense: God is temporal and exists at all times. As we see in Section 1.2.5, one prominent way of resisting the argument is by rejecting this assumption.

² In the foreknowledge literature this position is often called *theological compatibilism*. However, the label 'theological compatibilism' is also used to describe the view that God's determining activity is compatible with human freedom. Since we discuss that position at length in the next section, we use the term 'foreknowledge compatibilism' for the compatibilist view discussed in this section.

Omniscience is the property of knowing the truth value of all propositions. Essential omniscience is the property of necessarily knowing the truth value of all propositions. Infallibility is the property of being unable to make any mistakes in one's beliefs (Zagzebski 1996: 4–5). Infallibility is not entailed by omniscience (because an omniscient being might be capable of losing its omniscience), but it is entailed by essential omniscience.

The argument assumes that future contingent propositions can be true prior to the times they are about. Future contingent propositions are propositions about contingent future states of affairs. For example, the propositions *that you will decide to have cereal for breakfast on January 1, 2042*, and *that the Green Party will win a UK general election in 2027* are, at the time of writing (2018), future contingents. The future contingent propositions relevant to the argument are, of course, those which involve future human decisions.

Finally, the argument appeals to the intuitive idea that certain facts about the past are, to use Nelson Pike's words, "fully accomplished" and "over and done with" (1970: 59). These are what we might intuitively think of as 'genuine' facts about the past. In the literature on foreknowledge and free will they are called *hard facts* about the past. All other facts are called *soft facts*. The following facts are – relative to our time of writing in 2018 – paradigm cases of hard facts:

Paradigm hard facts relative to 2018

Martin Luther posted his ninety-five theses to the door of the castle church of Wittenberg in 1517.

Goran Ivanišević won Wimbledon in 2001.

It is very intuitive that hard facts such as these are now *fixed* in an important sense: no one has (now) the power to prevent these facts from being true; no one has (now) any choice about these matters. It is simply *too late* to do anything about them. Soft facts, by contrast, are not, relative to a given time, "over and done with", which is to say they are not, relative to that time, "temporally intrinsic" facts but are in part about other future times – they are "temporally relational" facts (see Todd 2013). As such, soft facts are not necessarily fixed in the sense given earlier. Suppose, for example, that it is now (in 2018, as we are writing this Element) true that you will choose to have cereal for breakfast on January 1, 2042. The fact that you will choose to have cereal for breakfast on January 1, 2042, is a soft fact because it is not (at the time of writing) over and done with. Moreover, this seems to be just the kind of thing that you do have control over: it's not too late for you to prevent your cereal eating. Not all soft facts which appear to be entirely (or mostly)

about the future are like this, however. For example, the fact that the sun will rise on January 1, 2042, is a soft fact which does not appear to be under anyone's control. So, some soft facts are fixed too. Then again, some facts about the past – or at least, partially about the past – seem to be soft facts which are not fixed. Thus, while *Goran Ivanišević's winning of Wimbledon in 2001* is a hard fact, relative to our time of writing, the fact – and let us suppose it is a fact – that *Goran Ivanišević's 2001 Wimbledon win is the only winning of the men's Wimbledon title by a Croatian in the twenty-first century* is, relative to the time of writing, a soft fact. The latter fact is in part about each Wimbledon championship that is to be played in the twenty-first century and no doubt many (present and future) Croatian tennis players hope they have the power to prevent its truth. Some soft facts which are in part about the past do appear to be beyond anyone's control, however; for example, the fact – assuming it is a fact – that Sally had breakfast yesterday, exactly two days before a meteor impact on Pluto, is a soft fact that appears to be beyond anyone's control. The distinction between hard and soft facts is discussed further in Section 1.2.3 which considers a response to the argument – Ockhamism – based on this distinction. Section 1.2.4 considers a response to the argument which challenges the idea that hard facts about the past are fixed.

When some proposition p is true, and some person S is unable at t to act so as to make it false, we say that S has no choice about the truth of proposition p . We use what has come to be a fairly standard abbreviation for this idea, namely, ' $N^S_t(p)$ '. Thus,

$N^S_t(p)$

is short for

p and S has, at and after t , no choice about the fact that p .

To illustrate with a concrete example,

$N^{\text{David}}_{2019}(\text{Goran Ivanišević won Wimbledon in 2001})$

is short for

Goran Ivanišević won Wimbledon in 2001 and David has, in and after 2019, no choice about the fact that Goran Ivanišević won Wimbledon in 2001.

That's quite a mouthful, hence the 'N' abbreviation.

With that background in place we can now present the argument for the conclusion that divine foreknowledge and human freedom are incompatible. We use Pike's example of Jones's lawn mowing as our ordinary action, although we cast it in terms of Jones's *decision* to mow her lawn, rather than the mowing

itself. Suppose, then, that it is true that at Saturday lunchtime (t_3) Jones will decide to mow her lawn. Time t_1 is some time before t_2 , which is itself a time before t_3 . The argument for the incompatibility of foreknowledge and free will runs as follows:³

- (1) God believed at time t_1 that Jones would decide at t_3 to mow her lawn.
- (2) $N^{\text{Jones}}_{t_2}(\text{God believed at } t_1 \text{ that Jones would decide at } t_3 \text{ to mow her lawn})$.
- (3) If $N^S_t(p)$ and $N^S_t(p \text{ entails } q)$, then $N^S_t(q)$.
- (4) $N^{\text{Jones}}_{t_2}(\text{God believed at } t_1 \text{ that Jones would decide at } t_3 \text{ to mow her lawn entails that Jones will decide at } t_3 \text{ to mow her lawn})$.
- (5) $N^{\text{Jones}}_{t_2}(\text{Jones will decide at } t_3 \text{ to mow her lawn})$.
- (6) If $N^{\text{Jones}}_{t_2}(\text{Jones will decide at } t_3 \text{ to mow her lawn})$, then Jones cannot decide to refrain from mowing her lawn.
- (7) If Jones cannot decide to refrain from mowing her lawn, she does not at t_3 decide to mow her lawn freely.
- (8) Therefore, Jones does not at t_3 decide to mow her lawn freely.

Clearly, there is nothing special about Jones deciding to mow her lawn, so the argument will generalise to all human decisions and actions. If sound, the argument establishes that divine foreknowledge of human decisions is incompatible with human freedom, a thesis we call *foreknowledge incompatibilism*. The version of the argument presented earlier yields a fatalistic conclusion. That is, it assumes that God is indeed omniscient and concludes – by establishing that foreknowledge and free will are incompatible – that free will does not exist. Few of those who defend the incompatibility of foreknowledge and free will, however, accept the fatalistic conclusion. Typically, they treat premise (1) as an assumption for an indirect proof and continue as follows:

- (9) But, Jones will decide at t_3 to mow her lawn freely.
- (10) Therefore, (contrary to (1)), God did not believe at t_1 that Jones would decide at t_3 to mow her lawn.

In contrast, *foreknowledge compatibilists* hold that divine foreknowledge and free will are compatible and so are committed to (1) and must therefore reject one or more of the argument's other assumptions, premises or inferences. All the responses considered in what follows are compatibilist except for the open theist response discussed in Section 1.2.6.

³ This argument is closely modelled on John Martin Fischer's (1989: 6) modal version of the argument.

1.2 Responses to the Argument

1.2.1 *Alternative Views of Free Will*

The argument for foreknowledge incompatibilism given earlier assumes that free will consists in having a choice about something. When an agent faces a *choice* or *decision* (we use the terms interchangeably), she has at least two options from among which she can select. Choice therefore entails the existence of what the contemporary literature calls *alternative possibilities*: possible alternative unfoldings of the world. (The existence of alternative possibilities does not entail choice, however, because there might be unfoldings of the world over which the agent herself has no control.) On this account of free will an agent, P, who faces a decision between A and B will decide freely only if (i) P is able to decide to A, and (ii) P is able to decide to B (where deciding to B might simply be deciding to refrain from A-ing).

This argument makes explicit use of this intuitive notion of free will in premise 7:

- (7) If Jones cannot decide to refrain from mowing her lawn, she does not at t3 decide to mow her lawn freely.

As such, one way of dispatching the argument quickly and effectively is to reject this view of free will. The most straightforward way of doing so is to say that the kind of control required for moral responsibility – this, recall, is the functional definition of ‘free will’ given in the introduction – does not involve choice. There are many such accounts of free will. Take, for example, Thomas Hobbes’s account. On his view an action is a behavioural event that is caused by a desire. And the agent herself is free in the execution of some action if there are no impediments to action that are extrinsic to the agent (Hobbes, Bramhall & Chappell 1999: 38). Having a choice about what one’s desires are, or which way to realise one’s desires, is simply not needed. Hobbes’s view does not fare well once it is accepted that a person may be subject to internal factors that she does not endorse but that determine or even merely influence what she does; and since Freud, it has been widely accepted that a person may have desires of which she is completely unaware and may not want to have.

Contemporary non-choice-based accounts of free will tend, therefore, to be more sophisticated. One example is Harry Frankfurt’s idea that freedom involves the alignment of a person’s ‘second-order’ desires with her ‘first-order’ desires. First-order desires are desires for food, shelter etc. Second-order desires are desires about one’s first-order desires. For example, to want to smoke is a first-order desire; to want to not want to smoke – i.e. to want to be free from any desire to smoke – is a second-order desire. The point is that you might want to

smoke while also wanting to not have that desire to smoke. Frankfurt's (1971) suggestion was that a person is free when a first-order desire which she wants to have – i.e. for which she has a second-order desire – causes her to act.

On non-choice-based views of free will such as Hobbes's and Frankfurt's, free will is straightforwardly compatible with God's foreknowledge. While such accounts solve the problem of foreknowledge, we say very little about them here, for two reasons. First, we remain unconvinced that any non-choice-based account of free will captures the notion of control relevant to free will. As we've already noted, choice has immediate and obvious moral significance; it is implicated in all or almost all areas of human social life and practice. And it seems to us that the degree to which non-choice-based accounts of free will appear plausible is the degree to which they smuggle back in the notion of choice. To give just one example, critics of Frankfurt's early hierarchical account pointed out that there doesn't seem to be anything special about second-order desires which makes them authoritative for the agent. Why think that an agent's will – the first-order desire which moves her to act – is free if she has a second-order desire for that first-order desire? After all, just as an agent might have first-order desires she would rather not have, so an agent might have second-order desires she would rather not have. This is a question of *identification*: which desires are truly the agent's own? (See Stump (1996) for further discussion.) In a later paper responding to this question, Frankfurt suggested that a desire might become authoritative for a person when the person *makes a decision* to identify with that desire: "Through his action in deciding, he is responsible for the fact that the desire has become his own in a way in which it was not unequivocally his own before" (1988: 170). This might well solve the problem, but if so, it is only because the key notion of *decision* or *choice* has been reintroduced.

The second reason for not treating non-choice-based accounts in any depth is that, even if someone presented a non-choice-based account of free will that sufficed for moral responsibility, we would still care to some degree about the freedom to do otherwise. That's because, as already stated, in our ordinary deliberation and the practices related to it we assume that we are free to do otherwise, and this assumption seems important to at least some of our judgements concerning that deliberation-based behaviour. And we would continue relying on this assumption *whether or not* our deliberation-based behaviours were considered things for which we could be morally responsible (Fischer 1989: 12).

1.2.2 Denying the Transfer Principle

Consider premise (3) from our argument for foreknowledge incompatibilism:

- (3) If $N^S_t(p)$ and $N^S_t(p)$ entails q , then $N^S_t(q)$.

Premise (3) is a version of what is known as the *Transfer of Powerlessness Principle*. Such principles are used in arguments for the incompatibility of free will with causal or natural determinism, as well as arguments for foreknowledge incompatibilism. Their purpose is to formalise the following intuition: if someone has no control over one fact, and that fact entails a further fact and the person has no control over that entailment, the person has no control over the fact which is entailed. To illustrate, suppose that it's currently Tuesday and Andy has no choice about whether it's Tuesday. Its being Tuesday today entails that it will be Wednesday tomorrow, and Andy doesn't seem to have any control over that entailment either. But then it's intuitive to conclude that Andy has no choice about the fact that it will be Wednesday tomorrow (see Fischer 1989: 7 for further examples).

The argument for foreknowledge incompatibilism appeals to the same idea: Jones has no choice about whether God believed at t_1 that she would decide at t_3 to mow her lawn (after all, t_1 could be a point in the distant past, e.g. 1 billion years ago); and she has no choice about the fact that God's past belief entails that she will indeed decide at t_3 to mow her lawn (because she has no control over God's infallibility); therefore, Jones has no choice about how she ends up deciding at t_3 .

Given how intuitive such examples are, the rejection of the Transfer of Powerlessness Principle will seem "extremely puzzling", as William Hasker has said, unless one can go beyond pointing out that the examples don't prove the principle to provide some positive argument for its rejection (2001: 102). Some philosophers have indeed attempted to do this. One of the most influential attempts was made by Thomas McKay and David Johnson (1996), who challenged a transfer principle that Peter van Inwagen used in an argument for the incompatibility of free will and causal determinism. Despite their article generating significant discussion, scholars have raised problems for it. One is that McKay and Johnson's argument employs a strong interpretation of the 'N' operator which requires the agent to be able to *ensure* that a given outcome occurs. But, as Timothy O'Connor (1993) has pointed out, arguments for incompatibilism do not need to employ such a strong reading of the 'N' operator; they need only require that the agent be able to do something which *might* have a particular result. This provides one way to escape the McKay and Johnson objection. Another problem for McKay and Johnson's argument is that it targets a specific formulation of the Transfer Principle. As such, their objection might be conceived of as a 'technical response' inasmuch as it exploits a technical flaw in the argument's formalisation, rather than addressing the incompatibilists' underlying worry. And because of this, many philosophers do not consider it decisive since it looks as if the Transfer Principle can be repaired