

Vihvelin and Fischer on ‘pre-decisional’ intervention

Simon Kittle
simon@kittle.co.uk
University of Sheffield

forthcoming in *Philosophia*

1. Introduction

A Frankfurt-style case (FSC) is a scenario in which an agent who performs an action, A, is intuitively morally responsible for A despite supposedly being unable to do anything different. The agent’s alternatives are apparently removed by the presence of an intervener who stands ready to intervene, should the agent be about to do other than A, thus ensuring that the agent A-s. Vihvelin is a prominent critic of such cases. She carves up Frankfurt-style cases up into two kinds: those employing *conditional intervention* and those employing *counterfactual intervention*. In cases of conditional intervention, the intervener is able to detect the very early stages of the victim’s action and intervene on that basis. For example, the intervener might detect when the victim makes a choice and intervene to stop that choice being acted on. Cases involving counterfactual intervention are those where the intervener is able to detect something *prior to* the agent’s choice, something that indicates how the agent would decide. Perhaps the intervener detects a twitch on the part of the victim (Frankfurt 1969, p. 835), an “inclination” to behave some way (Fischer 1994, p. 131), or some pattern of neural activity that is prior to (and not constitutive of) the victim’s decision. Whatever the details of the case, the key point is that the intervention is dependent on a *prior sign*: something that indicates how the victim will choose before the victim has made the choice. The distinction between conditional and counterfactual is not exclusive, inasmuch as the intervener might possess the powers needed for both kinds of intervention.

Vihvelin thinks that Frankfurt-style cases involving conditional intervention *couldn’t possibly* establish Frankfurt’s intended conclusion. This is because the agent always retains the ability to choose (or to “begin to choose,” whatever that might mean) otherwise. The focus, therefore, should

be on the remaining cases – those involving *counterfactual intervention*. These latter cases are contentious because they require the existence of a prior sign which reliably indicates how the victim will act. As such, the battle is usually over whether any such case is coherent. According to Vihvelin, both the proponents and the critics of these cases accept that *if* there is a consistent case of counterfactual intervention, *then* it will succeed in removing all of an agent's alternatives. Vihvelin considers this a mistake. She thinks that we can accept these cases as described by their proponents and still show that the victims in such cases have alternatives.

In this paper I enter a debate between Vihvelin (2000; 2008) and Fischer (2008) on this last point: if cases of counterfactual intervention are possible, do they succeed? To make her case Vihvelin presents a scenario purportedly parallel to a Frankfurt-style case where it is clear that the presence of a counterfactual intervener does not remove alternatives. Agreeing with Vihvelin's analysis of this case, Fischer has argued that it is not parallel to Frankfurt-style cases. And he constructs a hypothetical syllogism based argument to demonstrate that agents subject to counterfactual intervention cannot do otherwise. I argue that Fischer is wrong about the parallelism but right about cases Frankfurt-style cases, and I construct a hypothetical syllogism which is applicable to Vihvelin's parallel case.

Before that, a quick (but not incidental) note on terminology: Vihvelin's labels 'conditional intervention' and 'counterfactual intervention' are unhelpful. First, there is no difference in the modal status of the intervener or the intervention between cases of conditional and counterfactual intervention. In both kinds of case the intervener is actual, being as much a part of the scenario as the victim. And in both kinds of case the intervention is a non-actual possibility. The difference between the two kinds of intervention is a matter of *when* the intervention occurs in the so-called 'alternative sequence.' Second, the term 'counterfactual' is widely used in the literature on Frankfurt-style cases to describe any intervener who stands ready to intervene but whose intervention isn't needed (i.e. the standard usage covers both kinds of intervention). Furthermore, the term 'counterfactual' is unhelpful even when used in this standard way because it often leads to talk of the intervener (as opposed to the *intervention*) being counterfactual. This makes it easy to think that the intervener is benign – but this is to prejudge issues. As is now being recognised, the very presence of an intervener can and often does affect at least some of the agent's modal properties (capacities, capabilities,

abilities) (Levy 2008; Clarke 2011). As such, I will not use the term 'counterfactual' in either of these ways. I will use the label 'pre-decisional' to describe what Vihvelin calls counterfactual intervention and 'post-decisional' when referring to conditional intervention.

2. Agents subject to pre-decisional intervention are unable to do otherwise

Consider the following explicitly indeterministic Frankfurt-style case put forward by Vihvelin (2000, p. 20):

(COIN) At 6pm, Vihvelin and Player bet on a coin toss. Vihvelin bets heads, Player bets tails, and it comes up heads. The toss of the coin is an indeterministic process and is "genuinely chancy" with a 50/50 chance of coming up heads and tails. However, Vihvelin has a confederate called Black who is able to predict with "perfect accuracy" how the coin will land; Black also has the ability to act ahead of time to ensure the coin will come up how he wants. Vihvelin and Black prearrange that she'll bet heads and that Black will intervene, if needed, to ensure the coin comes up heads. Black knows, early in the morning, what the coin will do, and by noon any intervention that was needed is over and done with. At 12:01 Black "retires for the day."

This example is meant to parallel Frankfurt-style cases involving pre-decisional intervention. The conjunction of two facts might make it natural to think that the coin cannot come up tails here: first, Black's knowledge of what the coin will do is always right, second, Black is resolved to intervene whenever necessary. According to Vihvelin, however, to conclude on this basis that the coin cannot land tails would be a mistake. Vihvelin suggests that we need to recognise that there is a "complex truth" about the coin (2000, p. 18). On those occasions where, aside from intervention, the coin would come up heads, Black knows this and so doesn't intervene: in such cases the coin could have come up tails (because the coin toss was an indeterministic process). On those occasions where, aside from intervention, the coin would come up tails, Black knows this, intervenes and so the coin could not have come up tails. Vihvelin expresses this "complex truth" like so:

EITHER the coin comes up heads even though it could have come up tails OR the coin comes up heads and could not have come up tails (Vihvelin 2000, p. 18).

Vihvelin's point is that if the coin could come up tails (at least sometimes), then the agent in a pre-decisional Frankfurt-style case could have done otherwise. This is a significant conclusion because

most writers – both proponents and critics – have assumed that *if* there is a coherent case of pre-decisional intervention, *then* it will succeed in removing all of an agent's alternatives. That is why much of the literature on Frankfurt-style cases is about whether there is a valid case of pre-decisional intervention. On the one side, proponents of the Dilemma Defence are in effect arguing that cases of pre-decisional intervention are illegitimate (Widerker 2003); (Ginet 1996). On the other side, many proponents of Frankfurt-style cases attempt to develop nuanced cases of pre-decisional intervention – for example, the so-called buffer cases of Pereboom (2001) and Hunt (2005) – which aim get around the Dilemma Defence. They do this because they accept the conditional above: *if* there is a coherent case of pre-decisional intervention, *then* it will succeed in removing all of an agent's alternatives. Vihvelin thinks that much of this discussion is beside the point because even if there are cases of pre-decision intervention, the agent still has alternatives.

Fischer, who has made extensive use of Frankfurt-style cases, responded to Vihvelin by arguing that the analogy between her **COIN** case and the Frankfurt-style cases was problematic. Fischer accepts Vihvelin's assessment of the **COIN** scenario, but he puts this down to the fact that Black retires shortly after noon and so is not around when the coin is tossed (Fischer 2008, p. 335). This differs significantly from Frankfurt-style cases where the intervener stays on the scene to ensure that everything goes according to plan. Consider a typical Frankfurt-style case: Jones is the victim, Black the intervener wants Jones to A, and the question of whether Jones can refrain from A-ing arises. In standard FSCs, Black stays on the scene, ready and waiting to intervene, until Jones has finished A-ing. According to Fischer, this makes the following true:

If Jones were about to refrain from A-ing, he would not [succeed in refraining] (Fischer 2008, p. 335).

The truth of this counterfactual shows that Jones “lacks the power to choose and do otherwise” (Fischer 2008, p. 334). Fischer presents an argument to support his case, but before looking at it, it is worth introducing a second example. For one problem with **COIN** is its explicitly indeterministic nature. Although accepting the possibility of an indeterministic case concedes to the proponents of Frankfurt-style cases that which makes their case strongest, it also introduces complexities which tend to sidetrack the discussion.

To avoid these difficulties Vihvelin has developed an explicitly deterministic case. Such a case will not be accepted by incompatibilists, but Vihvelin's purpose in introducing it is to help us get clear about the different ways that pre-decisional and post-decisional intervention work. If, for sake of argument, we accept a context where pre-decisional intervention is obviously acceptable (i.e. a deterministic context), then we will see that Fischer's point about the intervener's effectiveness being dependent on him being around at the time of the victim's choice is incorrect. The following is a paraphrase of Vihvelin's case (Vihvelin 2008, pp. 354–355):

(BIKE) Jones lives in a deterministic universe. Jones is at home and has no plans for the evening. Sally phones him and asks whether he'd like to go for a walk or a bike ride. Jones thinks about it, decides he's not in the mood for a bike ride, but agrees to go for a walk. Later on, Sally pops round and off they go. Unbeknownst to Jones, Black, a Laplacean predictor, monitors everything he does. Black has "extensive and ultra-reliable knowledge" about the universe which permits him to predict everything that happens to Jones. Black's predictions "are always right." When Black predicts that Jones will, in the absence of intervention, do something contrary to Black's wishes, he hangs around and intervenes as necessary to ensure that Jones doesn't even begin to decide differently – Black renders Jones unable to do otherwise. This morning Black predicted that Jones would do just what Black wanted. As such, he retired for the day.

The deterministic nature of this example, combined with the distinction between pre- and post-decisional intervention, allows Vihvelin to answer Fischer's disanalogy complaint. First, we are to treat Black as solely a pre-decisional intervener (recall that he could be both). Black thus has the power to intervene *before* Jones even acts. How much before? Well, in a deterministic context with Black as a Laplacean predictor there is no reason why his prediction and intervention cannot be placed arbitrarily early. And once we can place Black's prediction and intervention at a large temporal distance from Jones' action, we can place Black himself at a large spatial distance from Jones at the time at which Jones acts. It's hard to see what grounds there are for resisting this move. Vihvelin thinks that Fischer's insistence that the intervener be on the scene at the time of action arises from not clearly distinguishing between pre- and post-decisional intervention. I think Vihvelin is right about this point. The case for the disanalogy between **COIN** and FSCs is faulty. Nevertheless, like Fischer, I want to defend the following claim:

If Jones were about to refrain, he would not [succeed in refraining] (Fischer 2008, p. 335).

Indeed, I think Fischer's defence (see below) of this claim is sound. I can endorse Fischer's defence despite disagreeing with him about the disanalogous nature of **COIN** and **BIKE** because, contra both Fischer and Vihvelin, I deny that the coin could have landed tails. I think there is an argument parallel to Fischer's which applies to the **COIN** scenario. If this is right, then the assumption typically made in the FSC literature, namely, that if there is a valid case of pre-decisional intervention, it would be a case where the agent has no alternatives, is vindicated.

Fischer's argument, which has the form of a hypothetical syllogism, is as follows (2008, pp. 338, 340):

(Hypothetical syllogism for a Frankfurt-style case)

- (1) If [the relevant agent] were about to refrain (in the absence of intervention by an external agent or factor), the triggering event would already have occurred.
- (2) If the triggering event had already occurred, Black would have intervened and forced Jones to act, in which case Jones would not have been able to refrain.
- (3) Therefore: If Jones were about to refrain, he would be rendered unable to refrain.

Vihvelin has pointed out that when the premises of this argument form are counterfactuals it is formally invalid. Fischer agrees, but points out that this doesn't mean there are no acceptable instances of the form, and of course, he thinks the argument above is an acceptable instance (Fischer 2008, p. 337). The general form of the argument above is as follows:

(Hypothetical syllogism involving counterfactuals)

- (1) If P had been the case, then Q would have been the case.
- (2) If Q had been the case, then R would have been the case.
- (3) Therefore: If P had been the case, then R would have been the case.

To see when instances of this formally invalid argument form are acceptable we need to look at how counterfactuals are assessed. Fischer summarises Lewis's semantics for counterfactuals as follows: a counterfactual of the form 'If P had been the case, Q would have been the case' is true (roughly) "just in case Q is true in the possible world or worlds in which P is true that is (or are) 'closest' ['most similar to'] to the actual world" (Fischer 2008, p. 338). What counts as *close* depends (in part) on the content of the counterfactual because in assessing counterfactuals we attempt to apply a principle of charity, and thus attempt to find a reading whereby the counterfactual comes out true. This opens up the possibility, when considering a hypothetical syllogism where both premises are counterfactuals, that the two premises "send us to different worlds" (in virtue of using a different similarity metric)

(Fischer 2008, p. 338). When this happens, the counterfactuals in the two premises are true in virtue of different possible worlds. Fischer describes this as “world-hopping” (Fischer 2008, p. 338). We find an example of this in a pair of statements discussed by Lewis (originally presented by Stalnaker):

- (1) If J. Edgar Hoover had been born a Russian, then he would have been a Communist.
- (2) If J. Edgar Hoover had been a Communist, he would have been a traitor.

Each of these has a reading where what is uttered seems true. The first might be naturally uttered if we were musing on how one’s place of birth affects what one believes; and it would be natural to take the second as true if we were discussing what sort of things render people traitors. The important point is this: to make each of these utterances come out true we have to appeal to different possible worlds. This is why if we attempt to use these statements in an argument, for example, by adding the following as a purported conclusion:

- (3) Therefore: If J. Edgar Hoover had been born a Russian, he would have been a traitor.

we will fail. More generally, it is the possibility of “world-hopping” that renders the counterfactual-involving hypothetical syllogism formally invalid. However, when there is a single world which grounds the truth of both premises, we do end up with an acceptable argument. Fischer contends that this is how it is with Frankfurt-style cases. Recall the hypothetical syllogism he endorses:

(Hypothetical syllogism for a Frankfurt-style case)

- (1) If [the relevant agent] were about to refrain (in the absence of intervention by an external agent or factor), the triggering event would already have occurred.
- (2) If the triggering event had already occurred, Black would have intervened and forced Jones to act, in which case Jones would not have been able to refrain.
- (3) Therefore: If Jones were about to refrain, he would be rendered unable to refrain.

Of this argument, Fischer says the following:

Given the story of the Frankfurt-type case, I do not see any reason to suppose that a structurally similar sort of world-hopping is taking place here. As far as I can see, the Frankfurt-story posits a single possible scenario in virtue of which the two premises are true. Thus, there is no reason to suppose that there is no possible world in virtue of which the conclusion is true; it is precisely the same single world in virtue of which the premises are true (Fischer 2008, p. 340).

With a couple of caveats, this seems right. The caveats concern potential ambiguities in the above argument. We have to understand 'about to refrain' not as referring to a time when Jones' refraining is imminent, but simply as a time when Jones, absent intervention, is set to refrain. In other words, 'Jones is about to refrain' pertains to some point after the prior sign. This secures the truth of (1). Premise (2) is problematic inasmuch as it potentially misses some cases: if the prior sign (triggering event) has occurred, Black may or may not have yet intervened. Black is free to intervene any time between the occurrence of the prior sign and some point just before the time of Jones's decision (how close this point is to Jones' action depends on how long his intervention takes). Thus, if the triggering event has occurred, then either Black will have intervened or he will be set to intervene (in good time). Premise (2) could be strengthened to reflect this as follows:

(2) If the triggering event had already occurred, Black would either (a) be set to intervene shortly, or (b) have already intervened, with the result that Jones will be or has been forced to act as Black wishes, in which case Jones would not have been able to refrain.

In a reply to Fischer, Vihvelin has argued that the underlying problem is not just that the hypothetical syllogism involves counterfactuals, but that it is what she calls a "back-tracking argument." As Vihvelin understands it, such arguments have the following form (Vihvelin 2008, pp. 344 n.4):

(The back-tracking argument form)

If the present were different in way D, then the past would have been different in way E

If the past were different in way E, then the future would be different in way F

Therefore, if the present were different in way D, the future would be different in way F

Such arguments reason from a counterfactual present to some counterfactual past which would permit or make possible that present, and then they reason forward again to a counterfactual future. Vihvelin says this reasoning is illegitimate when our interest is in "the causal upshots of some nonactual event or state of affairs" (2008, p. 357). She gives the following example to illustrate her point: suppose that we are wondering what would happen to sensible, cautious Sara, who currently has neither parachute nor hot air balloon, if she jumped off the roof. We would probably think that if Sara were to jump she would get hurt. But someone might come along and offer the following counter-argument (2013, p. 207):

Sara is cautious and so would jump off the roof only if she were securely strapped to a parachute or helium-filled balloon. As a result:

If Sara jumped, she would have been securely strapped to a parachute or a helium balloon.

If Sara had been securely strapped to a parachute or helium balloon, she wouldn't get hurt.

Therefore, if Sara jumped, she wouldn't get hurt.

Vihvelin thinks it obvious that this argument is bad, the problem being that it employs a back-tracking argument to answer a question about the causal upshots of some counterfactual event (Sara's jumping). According to Vihvelin, if "we evaluate these [causal] counterfactuals by considering possible worlds where the past is different – e.g. Sara is strapped into a parachute ... – we get the wrong results" (Vihvelin 2013, p. 208). However, the argument concerning Sara above is not faulty because it involves back-tracking. Rather, the problem once again is world-hopping. To see that the Sara story involves world-hopping note that we cannot consistently hold the following three statements:

(1) Sara would jump off the roof only if she were strapped to a parachute.

(2) Sara is not strapped to a parachute.

(3) Sara jumps off the roof.

To put it another way, if it is possible to ask what would happen were Sara to jump now (without a parachute), then (1) is false. This problem is nothing much more than an ambiguity: (1) is very plausible when read as saying something about what Sara would do *voluntarily*. And read in that sense the argument is fine. But then there is little sense in asking what would happen were Sara to voluntarily jump off the roof without a parachute. What Vihvelin wants to know about is the causal upshot of Sara's body being, say, a meter away from the roof's edge (in the wrong direction).¹ To make this point clearer, consider the following, which I contend is an acceptable back-tracking hypothetical syllogism (suppose that my wife and I each have our own set of keys but there are no spares):

(Acceptable back-tracking argument)

(1) If it were the case that my keys were not in my house, then it would've been the case that my wife took them with her by accident this morning

(2) If it were the case that my wife took my keys with her by accident this morning, then I would not have been able to leave my house this afternoon

¹ It is interesting to note that if Lowe's semantics of counterfactuals are correct then an ambiguity is exactly what we have here Lowe 1990, p. 81.

(3) Therefore: if my set of keys had not been in my house, I would not have been able to leave the house this afternoon

The premises here are causal counterfactuals in Vihvelin's sense and the argument exhibits the form of back-tracking reasoning. But the argument is acceptable, and this is because there is a single possible world in virtue of which the counterfactuals are true.

Now, Vihvelin puts into the mouth of the objector the following argument concerning **BIKE**:

(Vihvelin's backtracking hypothetical syllogism)

(1) If Jones had decided to ride his bike, Black would have known about it in advance.

(2) If Black had known in advance that Jones would decide to ride his bicycle, Black would have made it impossible for him to do so.

(3) Therefore, if Jones had decided to ride his bicycle, Black would have made it impossible for him to do so.

This argument, she contends, is faulty *because* it exhibits the back-tracking form. This is incorrect. As Vihvelin notes, we can only "agree that (3) is true, provided that we understand it as saying that the (relevant) closest worlds where Jones decides to ride his bike are all worlds where Black makes a *different* prediction in the morning" (Vihvelin 2008, p. 359). But 'different' here means 'wrong': Black makes a *wrong* decision. But of course, given the details of the case, Black *can't* make a wrong decision. That is impossible given the setup of the case. If we insist on asking what happens when Black makes a wrong prediction, then, whatever else we're doing, we're not talking about the same case. As with the Sara example, this argument is bad, not because it involves back-tracking, but because it involves world-hopping.

But the failure of this argument is no problem because Fischer's **Hypothetical syllogism for a Frankfurt-style case** is unaffected by the above point: it doesn't ask what would happen were Jones to *decide to ride his bike* (inconsistent with the details of the case) but instead asks what would happen if Jones *were about to* refrain (in the absence of intervention) from riding his bike. This secures the conclusion that the agent in a pre-decisional Frankfurt-style case could not do otherwise.

What then about Vihvelin's indeterministic **COIN** case? Vihvelin puts the following argument in the mouth of the objector:

(Vihvelin's hypothetical syllogism for COIN)

- (1) If the coin were about to land tails, Black would have predicted this and intervened
- (2) If Black had predicted this and intervened, the coin would be forced to land heads
- (3) So if the coin were about to land tails, it would be forced to land heads

Whether or not this argument is acceptable depends on how we read the phrase ‘the coin were about to land tails.’ One way of reading ‘about to land tails’ is as picking out an event which occurs immediately prior to the coin’s landing tails: the coin has been flipped, is currently falling through the air, and is ‘about to land tails.’ This would happen after Black’s prediction and intervention and such a reading requires us to suppose that Black made a wrong prediction and thus engage in world-hopping.

A second way of reading ‘the coin were about to land tails’ takes it to mean something like ‘the coin were about to land tails *in the absence of intervention.*’ Here the ‘about to’ carries no implication that anything will happen soon; rather, the idea is that things are set up such that, unless something changes, the coin would land tails at some point in the future. An objector might reply that, because the coin toss is an indeterministic process, there is no time at which ‘things are set up such that, unless something changes, the coin would land tails at some point in the future.’ But this is a worry about whether pre-decisional intervention is possible in the first place: can Black even be a perfect predictor of an indeterministic process? This objection is a problem for the proponent of Frankfurt-style cases. Vihvelin’s project is that of attempting to ask, on the assumption that pre-decisional intervention is possible, whether it succeeds. If pre-decisional intervention is not possible, then Vihvelin’s argument (and my reply) is perhaps moot, but the objection is not itself an objection to Vihvelin’s points nor to my reply.

The addition of the clause in italics strengthens the argument because it ensures that all cases are covered: there are those cases where (absent intervention) it’s going to land heads and there are those where (absent intervention) it’ll land tails – we assume for simplicity that the coin cannot land on its side or spontaneously combust. But that means we have all bases covered: the coin cannot land tails. The suitably amended version of the argument therefore shows that the following is false:

If the coin were about to land tails (in the absence of intervention), no outside force would make it land heads.

As the truth of this statement is a necessary condition on the coin's being able to land tails Vihvelin is wrong and her "complex truth" is instead a complex falsity.

What, in closing, should we say about Vihvelin's intuition pumping with respect to the **COIN** case? It is certainly possible, especially with Vihvelin's helping hand, to feel the pull of the intuition that the coin could come up tails: after all, Black has all his work done by noon such that by the time of the coin toss he's fast asleep (Vihvelin 2000, p. 16). How could we possibly doubt that the coin could land tails given that the coin toss is an indeterministic process and the coin's environment fair? The answer, I think, is related to the objection mentioned above concerning the very possibility of pre-decisional intervention. Many have the intuition that pre-decisional intervention in an indeterministic context just isn't possible: Black cannot have the powers and abilities to perfectly predict the outcome of an indeterministic process such as the coin toss is stipulated to be. Vihvelin is sympathetic to these worries (Vihvelin 2000, p. 18). But she wants to bracket these intuitions and nevertheless ask whether the intervention, if it were possible, would be successful. But if we are bracketing our intuitions concerning the possibility of the scenario as a whole, then we should bracket our intuitions concerning what it is possible for individual entities in that scenario to behave. So the intuition pumping is no threat to the conclusion of my argument.

3. Conclusion

Vihvelin and Fischer disagree over whether there is a successful hypothetical syllogism based argument which shows that agents in (pre-decisional) Frankfurt-style cases cannot do otherwise. Vihvelin does not think there is, and aimed to show this with her **COIN** case. Fischer agreed about the **COIN** case but argued that it was disanalogous to Frankfurt-style cases. I agreed with Vihvelin that the cases are analogous, but contrary to both Fischer and Vihvelin I argued that there that there is a hypothetical syllogism available for each of **COIN** and **BIKE**. My argument here relied on resources in Fischer's argument about 'world-hopping.' The result is that *if* pre-decisional intervention is possible, agents subject to such intervention cannot do otherwise. This result holds for both deterministic and indeterministic Frankfurt-style cases. The discussion, therefore, over whether there is a valid case of pre-decisional intervention is not (as Vihvelin thinks) beside the point. The Dilemma Defence, and the Frankfurt-style cases developed in response to it, are not redundant.

Publication bibliography

- Clarke, Randolph (2011): Responsibility, Mechanisms, and Capacities. In *The Modern Schoolman* 88 (1), pp. 161–169.
- Fischer, John Martin (1994): *The Metaphysics of Free Will: An Essay on Control*: Wiley-Blackwell.
- Fischer, John Martin (2008): Freedom, Foreknowledge, and Frankfurt: A Reply to Vihvelin. In *Canadian Journal of Philosophy* 38 (3), p. 327.
- Frankfurt, Harry G. (1969): Alternate Possibilities and Moral Responsibility. In *Journal of Philosophy* 66 (3), pp. 829-39.
- Ginet, Carl (1996): In defense of the principle of alternative possibilities: Why I don't find Frankfurt's argument convincing. In *Nous* 30, pp. 403–417.
- Hunt, David P. (2005): Moral responsibility and buffered alternatives. In *Free will and moral responsibility*, p. 126.
- Levy, Neil (2008): Counterfactual intervention and agents' capacities. In *The Journal of philosophy* 105 (5), pp. 223–239. Available online at <http://cat.inist.fr/?aModele=afficheN&cpsidt=20899757>.
- Lowe, E. J. (1990): Conditionals, context, and transitivity. In *Analysis* 50 (2), pp. 80–87.
- Pereboom, Derk (2001): *Living without free will*: Cambridge Univ Press.
- Vihvelin, Kadri (2000): Freedom, foreknowledge, and the principle of alternate possibilities. In *Canadian Journal of Philosophy* 30 (1), pp. 1–23.
- Vihvelin, Kadri (2008): Foreknowledge, Frankfurt, and Ability to Do Otherwise: A Reply to Fischer. In *Canadian Journal of Philosophy* 38 (3), p. 343.
- Vihvelin, Kadri (2013): *Causes, laws, and free will. Why determinism doesn't matter*. New York: Oxford University Press.
- Widerker, David (2003): Libertarianism and Frankfurt's Attack on the Principle of Alternative Possibilities. In Gary Watson (Ed.): *Free Will*. 2nd ed.: OUP Oxford.