

Ethics // Spring 2026

Handout 7

Coercion: Nozick

PRELIMS. The task is to explicate coercion. We are after its necessary and sufficient conditions. But why coercion? Because it connects to freedom:

Coercion is not necessary for unfreedom: The person unaware that he will be caught and punished if he robs the bank is unfree to rob it. But he is not coerced.

Coercion is not sufficient for unfreedom: A successful bluff coerces you into not acting, yet you were not *in fact* unfree to act (the bluffer had no means to follow through).

Thus the target:

(7-1) P coerces Q iff Γ .

Question 1. Do you agree with Nozick's initial intuitions about the coercion-freedom relation?

UNDERSTANDING THREATS. In a highly recommended book, Hart and Honoré's proposal for Γ is:

441

(7-2) Γ : (1) & (2) & (3) & (4) & (5).

We have the following:

(7-3) P coerces Q into not doing A iff:

- (1) P threatens to bring about some consequence if Q does A (and knows it).
- (2) The threat makes doing A substantially less eligible as a course of conduct for Q .
- (3) P makes the threat in order to get Q not to do A , intending Q to recognise the threat.
- (4) Q does not do A .
- (5) P 's words or deeds are part of Q 's reason for not doing A .
- (6) Q knows that P has threatened to bring about the consequence if Q does A .

The key notion in Γ , as can be seen from (1) and (3), is threat. But the concept of threat is itself puzzling! So we have two problems right from the start: we are not sure what counts as a coercion and how it relates to a threat, but we are also not sure what counts as a threat to begin with.

Thus suppose I'm telling you:

441

(7-4) One more step, and you are dead!

I intend to coerce you ('make you') not to move by threatening you. You oblige: you don't want to die. However, it may be that you understood my utterance as a warning. For example, you saw a rock hanging dangerously above, whereas I meant that I would shoot you. So the conditions (1)-(5) seem to be satisfied, but you were not *really* coerced. I thought I was coercing you, but in fact you were just being cautious. We need a condition to refer to the 'uptake' on the part of the speaker. Hence Nozick introduces an additional condition (6): the hearer must know that the speaker has threatened him.

Remark 2 (Uptake). Nozick is not entirely right on this issue. The problem has to do less with the concepts of coercion or threat and more with the idea of a speech act. Threats are a kind of a speech act. Now for a speech act to be successful ('felicitous') the hearer must have some basic understanding of what's going on. Suppose there is a raging bull about to charge, and I am shouting to you, 'Watch out!' But you are distracted: either you haven't heard me at all, or didn't understand my English, or didn't think it was a warning. It is plausible to say: I *tried* to warn you, but failed. Or suppose that I am chairing a meeting and saying to the audience, 'I thereby declare the session open.' But no-one has heard me. Once again, I tried to open the session, but failed. Thus the speech act must, for it to *be* a particular act, secure the hearer's 'uptake': the hearer must take the utterance in a certain way, namely, to identify it *as* a warning, a command, a declaration. Therefore, Hart and Honoré who were well aware of Austin's theory of speech acts, would probably regard (6) as superfluous: it is built already into the condition (1). See Austin, *How to Do Things with Words* and Strawson, 'Intention and convention in speech acts'.

The next problem is this. Suppose I am shouting (7-4) to you. You do understand this as a threat. You also understand that I don't want you to move. You understand that this is why I have threatened you. However, you don't want to move anyway (you are tired). Then you realise how upset I will be, and you don't want to upset me. So you move for *that* reason.

Here, then, there is a threat, and the hearer understands it, but he is not complying with the threat *because* of it. That presumably is why we don't think of him as being coerced. We need to change (5) to (5'). We need to make explicit the reasons that motivate the hearer to comply with the threat. For coercion, the motive is a threat of some painful (evil) consequence:

(5') Part of *Q*'s reason for not doing *A* is to avoid or lessen the likelihood of the consequence *P* has threatened to bring about.

Correspondingly, we also we need to replace (1) and (2) with:

(1') *P* threatens to bring about or have brought about some consequence if *Q* does *A* (and knows he is threatening to do this).

(2') With this threatened consequence, *A* is rendered substantially less eligible as a course of conduct for *Q* than it was without the threatened consequence.

The key shift is the phrase 'threatened consequence'. The original (2) just said the threat makes doing *A* less eligible. This leaves open that *Q*'s deliberation is affected by something other than the consequence *P* is specifically threatening. (2') closes that by tying the eligibility reduction specifically to *that* consequence, the one *P* is threatening to bring about.

THE SSRC EXAMPLE. Should we accept (3)? That is, should we always wish to bring about the action that we are coercing the hearer to perform? Normally yes. But not always. Nozick's example involving SSRC is an especially weird one. Let me split it into two.

In one situation, perhaps in an ordinary science experiment, I don't want you to comply with the threat—I'm just assembling the data. But then, I never intend to follow up on my threat. Then this is a mock threat, not a real one. Suppose you comply with the threat—or rather, you comply with what *you* take to be a threat. I shout (7-4), though not intending to do anything to you. You think I am very serious, and you stand still. Were you coerced to stand still? Can you be *really* coerced by *unreal* threats?

I am not sure what to say. On the face of it, to say that you can sounds strange: you were coerced, but nobody was (really) coercing you. I wonder what the legal theory says. You, for example, may use my utterance (a mock threat) as an *excuse* for standing still and not saving someone, say. But I can also plead not guilty: 'I was joking, it wasn't a threat, you see.'

Nozick's point, I think, is in any case unrelated to mock threats: the SSRC researchers are allowed to kill the subject! Let's have a more realistic example: Suppose I threaten to kill you if you don't give me the money. But at the same time, I made a bet: if you comply, I lose the bet, and if you don't, then I win. My interest is that you do *not* comply. However, if you don't comply, I will still follow up on my threat and kill you (at least I intend to do so). So it is a regular threat.

Then we have: if you comply with the threat, you would be said to be coerced, yet my overall interest is your non-compliance. So (3) fails. So (3) is not necessary for coercion. Instead of (3), we must say something quite different, namely: part of my reason for deciding to kill you is that this would make your not giving the money worse (in your eyes) than giving the money.

Remark 3 (Non-threatening warnings). As Nozick notes later, this condition distinguishes threats from warnings. Briefly, the consequences indicated in the warning may worsen your non-compliance. But unlike threats, that is not my motive for issuing the warning.

Finally, there is a situation where a threat (if it is a threat) doesn't allude to an evil to be suffered as a result of ignoring it. We consider a scenario where you want to ϕ in order to achieve a goal *g*, but I threaten you with doing something that will simply prevent you from achieving *g*. Nozick's example is this:

(7-5) If you say another word, I'll turn my hearing aid off.

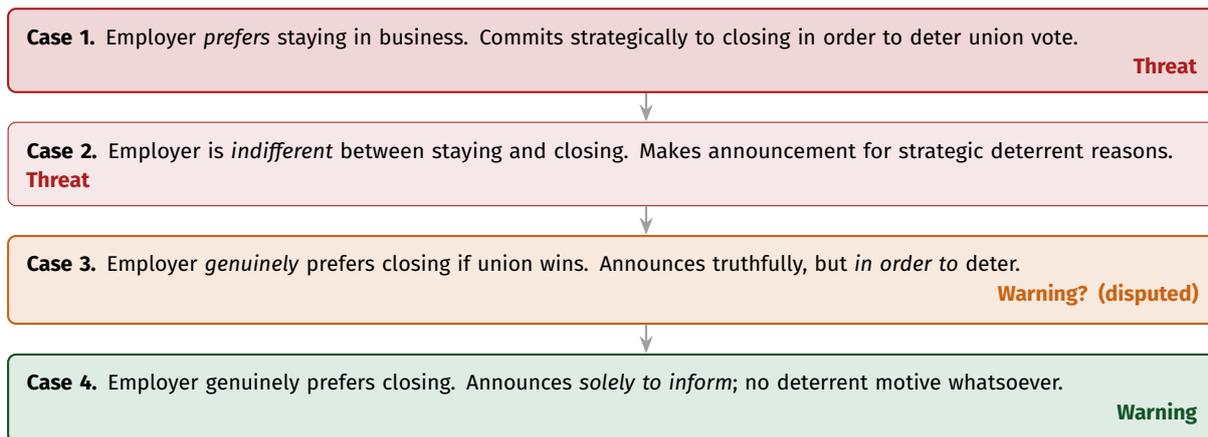
Here, the action is 'speaking' and the goal is 'having the audience hear/understand what you say'. Suppose you want me to hear what you say, but I'm announcing (7-5) to you. Nozick reasons that this is not the case of coercion: I didn't coerce you to shut up. As I understand him, intuitively this is because you wouldn't have suffered by disobeying my threat, though you *were* motivated to obey it. Hence the condition (7) where the key word is 'worse off'.

Question 4. Show that cases (1)-(3) of various mishaps fail the criteria of coercion. Should the case of a mock threat be analysed similarly?

DETERRENCE. Nozick then briefly considers the case of a ‘deterrence’. Here, I may arrange things in such a way that it would be obvious to you that bad consequences will follow should you decide to ϕ . In contrast to the standard case, there is no threat: I am not *telling* you, in some way, not to ϕ —although perhaps I *am* telling you what consequences will follow. It seems to me that a lot here will depend on how we think about arranging those bad consequences. Did I arrange them specifically to prevent you from ϕ -ing, or did I do that for some other reason? In the latter case, you can’t say:

(7-6) He forced me not to ϕ !

I could object that *I* didn’t force you—that is, it wasn’t my intention, it wasn’t what I was doing. The circumstances simply occur that prevented you from ϕ -ing. In the former case, when I arrange things overtly with the intent to prevent ϕ -ing, it seems that you can say something like (7-6). There are other related situations to consider here, too. See Figure 1.



Key variable: Is ‘worsening *Q*’s alternative of electing the union’ *part of P*’s reason for making the announcement? Cases 1–2: Yes (strategic deterrence). Case 4: No (pure information). Case 3: Action would occur anyway, announcement is truthful—but deterrence is still a motive.

Figure 1: Threats and warnings

Case 3 is the hard one. Consider:

Example 5 (Teacher). The teacher who warns a student they will fail, in order to motivate them to work harder, is plausibly *warning* rather than *threatening*—even though the motive is to change behaviour. Nozick’s intuition is that since the teacher does not *decide to fail the student in order to worsen their alternative of being lazy* (the failing would happen anyway), there is no threat. The deterrent intent attaches only to the announcement, not to the underlying action. Critics may push back: if the announcement is made *in order to* influence behaviour, the distinction from threats may be unstable.

PERIPHERAL CASES. To be omitted.

THREATS AND OFFERS. We now look at the difference between threats and offers. On the face of it, when I offer you a reward, I may be said to coerce you to accept it: if you refuse, you would be worse off, just like when you refuse to comply with threats.

Nozick suggests to distinguish between threats and offers by appealing to the ‘normal course of events’ (usually called ‘baseline’ in the literature):

- (7-7) a. *P* threatens *Q* not to ϕ : *P* says that, following *Q*’s non-compliance and ϕ -ing, he will bring about a consequence that is, all things considered, *worse* than a *normal* consequence of ϕ -ing.
- b. *P* offers *Q* not to ϕ : *P* says that, following *Q*’s non-compliance and ϕ -ing, he will bring about a consequence that is, all things considered, *better* than a *normal* consequence of ϕ -ing.

This is illustrated in Figure 2.

A relatively minor problem is how to judge ‘better’ and ‘worse’. A more serious problem is how to identify the ‘baseline’.

This is seen in Nozick’s examples of drug dealing in (a) and (b) in the text. In (a) we are supposed to say that *Q* was threatened with withholding, hence coerced, whilst in (b) he was offered a drug in exchange for beating up a person, hence not coerced.

I agree about (b), but is (a) right? Wouldn’t it make more sense for *P* to say in (a), ‘The terms of my offer have changed’? Suppose I’ve been buying bread at the shop for 10TL for the past year. One day the shop raises

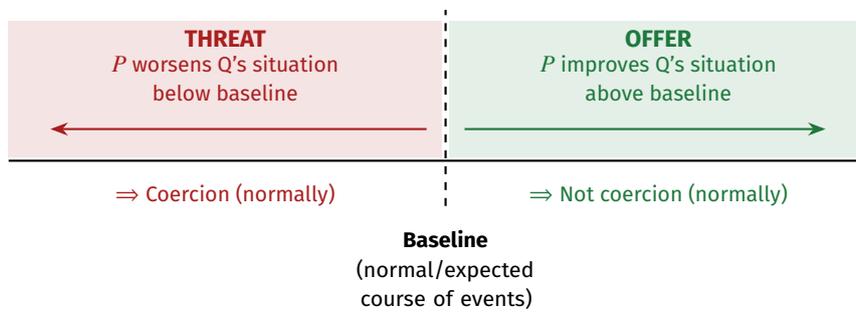


Figure 2: Threats and offers

prices (well above the inflation rate, too). Maybe that's a bad offer, an unfair offer, but it's not a threat, surely. One wishes to say: the utterance in (b) is made in the framework of offers already in place. The thorny issue is, therefore, how we think of the 'normal course of events'.

Remark 6. Nozick says that in (a) *P* made both a threat and an offer. But the reason he made an offer (so Nozick) is that, normally, people don't get drugs for beating others. That's not the baseline that seems plausible to me. 448

This problem of understanding the baseline is obvious in the example of saving a drowning man. If saving men is a normal routine, then the rescuer is threatening the man. If letting people drown unassisted is a normal routine, then the rescuer is making an offer. 449f

Perhaps normativity should play a role here. What matters is not a statistically regular routine, but a normatively regular one. This is evident in the slave example. Since beating up slaves—indeed, owning slaves in the first place—is not morally acceptable, it can't count as a baseline. 450

This intuition to me seems weak. I'd rather talk of 'immoral' offers and choose a morally neutral baseline.

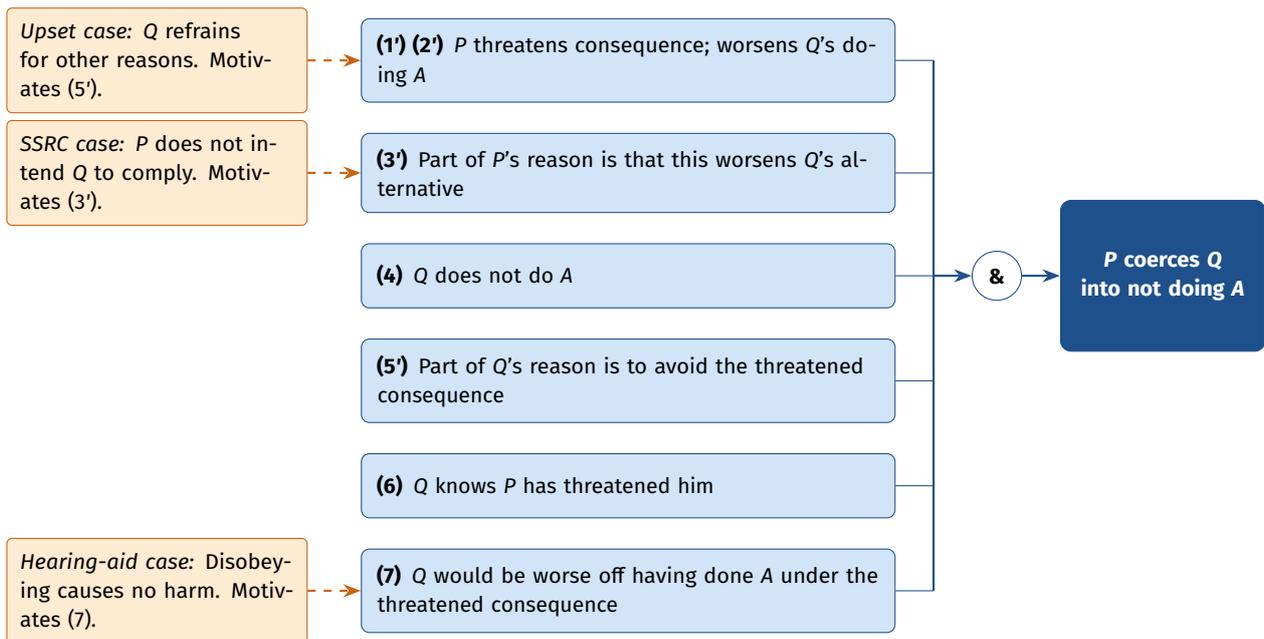


Figure 3: Definition of coercion

YSB