Article:

The list is long of those who, in the history of ethics, have advocated the thesis that what one ought to do is the best that one can do. This thesis has had its detractors, of course. I wish, however, not to defend it against external attacks but to adjudicate a dispute internal to it. The dispute concerns what have come to be called actualism and possibilism. In the first section, I shall argue that actualism has a number of unacceptable implications; in the second section, I shall defend possibilism against the charge that it has a number of unacceptable implica-tions. In this way, I shall attempt to show the superiority of possibilism over actualism. Then, in the final section, I shall elaborate on possibilism's account of where it is that an agent goes wrong when he does go wrong.

Part I

Actualism is the view that whether or not an action (or course of action) is the best that an agent can do is a function of what would happen if he were to perform that action (or course of action). The most straightforward formulation of this view is this:

> An agent S ought to do an act X if and only if what would happen if S did X is better than what would happen if S performed any alternative action.

Possibilism, in contrast, may be formulated as follows:

> An agent S ought to do an act X if and only if what could happen if S did X is better than what could happen if S performed any alternative action,

where "could" expresses personal possibility, and so "what could happen" means what it is possible for S to have or make happen.[1] In many cases, these two doctrines will yield the same account of what an agent ought to do. In fact, the distinction between actualism and possibilism was first brought to light, so far as I know, only very recently. At first, actualism was favored;[2] then possibilism was touted;[3] actualism then received a vigorous defense.[4] It is possibilism's turn once again.

Actualism and possibilism will often yield different accounts of what an agent ought to do when his own future actions figure in the evaluation of his alternatives. A by-now-standard sort of case where this occurs is the following. Jones needs comments on a paper soon and asks me for some. It would be best if I accepted the task (on Monday, say) and provided the comments on time (on Wednesday, say). It would be worst if I accepted the task and failed to provide the comments on time. In between, in terms of

value, lies the option of not accepting the task (and, therefore, of not providing the comments on Wednesday). Unfortunately, if I were to accept the task on Monday, I would not provide the comments on Wednesday, although I could do so. Ought I to accept the task?

In this case, my options are these (ranked in terms of value):

    1: A(ccept) & C(omment)
    2: —A & —C
    3: A & —C.

According to possibilism, I ought to accept; according to actualism, I ought not to accept. I think possibilism gives the correct answer here.
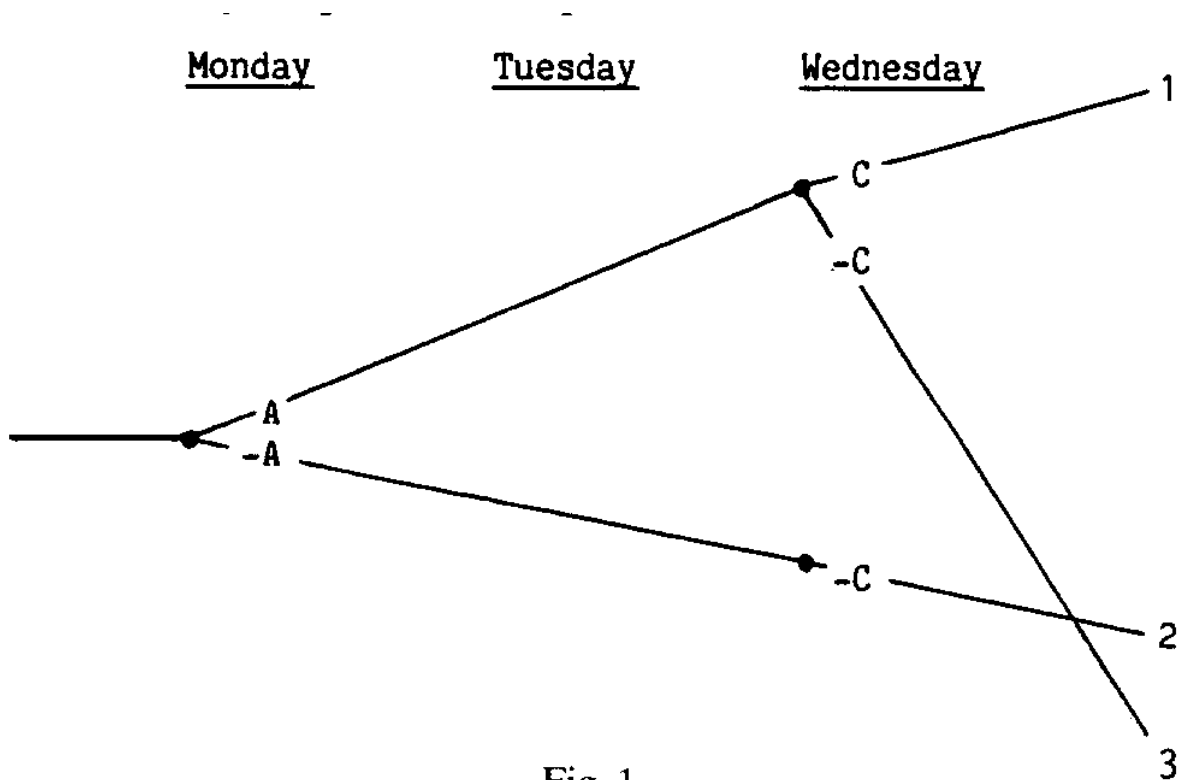
My situation may be pictured in part as follows:[5]

Fig. 1.

This sort of picture has its limitations, but it is nonetheless helpful and, I think, suggestive. It illustrates how my options on Monday are to accept and not to accept. (What I shall call Monday's "point of choice" is represented by the black dot.) If I accept, I have the options on Wednesday of providing comments and not providing comments. If I do not accept, my only option on Wednesday is not to provide comments. (Here the dot does not represent a point of choice, since there is only one line leading from it.) The three courses of action are ranked according to their relative values.

Suppose that I take the third-ranked course of action. Where did I go wrong? According to actualism, I went wrong at both Monday's and at Wednesday's points of choice; according to possibilism, only at

Wednesday's. The picture is intended to express the intuitive appeal of possibilism. By accepting the task, I set off on the right track; on Tuesday, I was still on the right track; it was only on Wednesday that I took a turn for the worse, that I veered from the straight and narrow, that I lapsed morally.

Part I.i
It might seem that possibilism is obviously preferable to actualism. Compare it with actualism in another version of this same case, one where I take the second-ranked course of action. Insofar as this course is not ranked first, it is clear that, according to the thesis that what one ought to do is the best that one can do, I went wrong somewhere. According to possibilism, of course, I went wrong at Monday's point of choice. But where, according to actualism, did I go wrong? Nowhere, it seems. Track 2 is the right track, after all; had I accepted the task, I would have wound up on Track 3. But isn't it plainly false to suggest that I didn't go wrong at all?
The actualist has a response here. He will point out that, according to him, by not accepting the task of commenting on Jones's paper I did do wrong. I didn't do wrong with respect to not-accepting, but I did do wrong with respect to not-accepting-and-not-commenting. Why? Because, given the rankings indicated in Figure 1, we may note that, even though it is true that doing A would have turned out worse than its alternative, it is nonetheless true that doing both A and C would have turned out better than any of its alternatives.

Is this an acceptable response, though? I think not, for it commits the actualist to a denial of several very plausible theses, theses to which we often appear to appeal in our moral reasoning.

First, the actualist must deny the following:

> (T1) If S ought to do both X and Y, then S ought to do X and S ought to do Y.

This is demonstrated, in the present case, by the fact that doing both A and C is obligatory while doing A is not. Some actualists have explicitly and boldly embraced this implication. Frank Jackson and Robert Pargetter, the authors of the recent, vigorous defense of actualism, say this:

> Perhaps an overweight Smith ought to stop smoking and eat less, but it may not be true that he ought to stop smoking. For it may be that were he to stop smoking he would compensate by eating more.[6]

But imagine Smith's quandary. He's been told that he must stop smoking and eat less; he therefore resolves to stop smoking, only to be told, "Oh no, you mustn't do that!"

Notice that (T1) is implied by the following:

> (T2) If S ought to do X, and it is logically impossible that he do X without doing Y, then S ought to do Y.[7]

Thus actualism implies that (T2), also, is false. Some may be heartened by this, since they take it that some version of the Good Samaritan Paradox shows (T2) to be false anyway. Not so. Properly construed, (T2) is not touched by this paradox.[8] Moreover, questions about the paradox aside, (T2) is surely very plausible; why else call the paradox a paradox?

Notice, further, that (T2) is itself implied by the following:

(T3) If S ought to do X, and he cannot do X without doing Y, then S ought to do Y.

(Here "cannot" has to do not with logical possibility but with personal possibility.)[9] Thus the actualist must deny (T3), too, despite what would appear to be its great plausibility.

The actualist must also deny this thesis:

(T4) If S ought to do X, and he ought to do Y, then S ought to do both X and Y.

Let X be my accepting-and-commenting and Y be my not-accepting. I cannot both accept-and-comment and not-comment; hence, according to the thesis that I ought to do the best I can, it is not the case that I ought both to accept-and-comment and not-comment. Yet each, according to the actualist, is such that I ought to do it.

Similarly, the actualist cannot accept this:

(T5) If S ought to do X, and he ought to do Y, then S can do both X and Y.

For (he says), I ought to accept-and-comment and I ought to not-comment; yet I cannot do both. Now, some would deny (T5) anyway, claiming that moral dilemmas show it to be false. While I reject this line of argument,[10] the point to be made here is that the actualist does not reject (T5) because he accepts the possibility of moral dilemmas. Suppose that what would prompt me not to comment on Jones's paper, were I to accept the task of doing so, is sheer bloody-mindedness; I just "couldn't" (more accurately: wouldn't) be bothered. Even in this case, the actualist would say that ought to not-accept. Yet he would also say that I ought to accept-and-comment. Now, I cannot do both, but it would surely be a gross distortion of my situation to portray it as a moral dilemma; for the actualist must admit that I could do all that I ought (namely, accept-and-comment, the doing of which would imply that it is not the case that I ought to not-accept).[11]

This brings out what is perhaps oddest about the actualist's view, and that is that it sanctions wholly avoidable wrongdoing — surely some-thing that a theory of wrongdoing ought itself to avoid. Why should the bloody-mindedness that I display in not commenting on Jones's paper be accommodated in this way? Why should my moral obligations be thought to be thus tailored to my own easily avoidable moral failings? Of course, my failure to comment might (in some other version of the case) be due not to bloody-mindedness but rather to a wholly understandable falling-short of some goal which, though possible for me, would have been especially difficult to attain. And it might seem that what actualism says in this version of the case is quite plausible. I could (just) get the comments done in time, but I'm not perfect and, despite my good intentions when accepting the task, I, like almost all others placed in such a situation, simply wouldn't get them done in time; surely it is only being fair (or realistic, or practical) to say that under the circumstances I ought not to accept the task. Yet, whatever plausibility there is to this position (I do not find it persuasive; it seems to me to trade on the failure to distinguish wrongdoing from blame-worthiness, a point to which I shall return at the very end of this paper), the actualist can find no solace in it. For the position seems to presuppose the view that one ought to do the best that can reasonably be expected of one (or something like that), rather than the view that the

actualist avowedly advocates, namely, that one simply ought to do the best one can. There is no escaping the fact that actualism sanctions avoidable wrongdoing (it implies that, on occasion, S ought to do X, even though his doing X requires that he do wrong and even though he can avoid wrongdoing altogether), whether or not such wrongdoing reflects bloody-mindedness or some other morally defective character-trait.

Part I.ii

As Jackson and Pargetter point out, the diachronic nature of the commenting case is not essential to the discussion. The following case shows why.[12] Suppose that I am driving through a tunnel behind a slow-moving truck. I may either change lanes or not, and either accelerate or not. My options are to be ranked as follows:

> 1: —C(hange)   & —A(ccelerate) (simultaneously)
> 2: C & A
> 3: C & —A
> 4: —C & A.

Suppose that, whether or not I accelerate, I am going to change lanes. Then actualism implies that it is not the case that I ought to not-accelerate, even though it is the case that I ought to not-change-and-not-accelerate. (Possibilism, of course, implies that I ought to not-accelerate.) This case thus brings out, just as well as the commenting case, the fact that actualism implies the falsity of (T1) through (T5).

Jackson and Pargetter themselves subscribe to a version of actualism slightly but importantly different from that with which we have been working so far. It is this:[13]

> S ought to do X if and only if what would happen if S did X is better than what would happen if S did not do X.

This version is no more successful than the first. It, too, implies that (T1) through (T5) are false; for it yields the same prescriptions as the first version in both the commenting and the driving cases. Indeed, Jackson and Pargetter's version seems to me even more problematic than the first, for it yields some counterintuitive prescriptions that the first does not. Consider this case, which Jackson and Pargetter them-selves provide.[14] Tom, Dick, and Harry are standing for election. My voting options are ranked thus:

> 1: (vote for) T(om) & —D(ick) & —H(arry)
> 2: —T & D & —H
> 3: —T & —D & H
> 4: —T 8c —D & —H.

Given this ranking, we may suppose that the following is true: if I were to vote for Tom, that would be better than if I were to perform any alternative action (that is, better than if I were to vote for Dick, or Harry, or no one); thus it is not the case that, if I were to vote for Dick, or Harry, or no one, that would be better than if I were to perform any alternative action (for it wouldn't be better than my voting for Tom). But note this. It is consistent with the foregoing ranking to suppose not only that, if I were to vote for Tom, that would be better than if I did not vote for Tom, but also that, if I were to vote for Dick, that

would be better than if I did not vote for Dick; for we may suppose that, if I didn't vote for Dick, I'd vote for Harry. Hence, while both versions of actualism imply that I ought to vote for Tom, the second (but not the first) version implies also that I ought to vote for Dick. This implication is surely to be rejected.

Part II
How successful are the foregoing criticisms of actualism? This is debatable. I argue: If P then Q (if actualism is true, then such-and-such a thesis is false), not- Q,15 hence not-P. Adherents to actualism may well respond: If P then Q, P, hence Q. This is not an uncommon situation in philosophy, and in this case, as in others, an indisputable resolution of it is probably not forthcoming. Nonetheless, I shall now try to strengthen the possibilist's position by these means: rebutting objections raised against possibilism by actualism's most prominent proponents, Jackson and Pargetter; pointing up an inconsistency in Jackson's and Pargetter's own position; and showing how possibilism can accommodate much, if not all, of what allegedly makes actualism so plausible.

Jackson and Pargetter give three objections to possibilism. I shall rebut each in turn.

Part II.i
The first objection is this. Detachment for conditional obligation is valid. That is, the unconditional obligation to perform an action may be inferred from the conditional obligation to perform it when the condi-tion in question is satisfied. Thus, given both the fact that I ought to decline if I won't provide the comments in time and the fact that I won't provide the comments in time, it follows that I ought to decline (period).[16]

The fact is that detachment for conditional obligation is simply not (unrestrictedly) valid. I ought to vote for Tom, but I won't; if I don't vote for Tom, I ought to vote for Dick; it hardly follows that I ought to vote for Dick. Elsewhere, Jackson and Pargetter recognize this very point.[17] But then they try to save their original thesis by drawing a distinction between types of conditional obligation. First there is what they call hypothetical obligation. A statement of hypothetical obligation has the form "If it were the case that p, S would be obligated to do X." As Jackson and Pargetter note, detachment with respect to such a subjunctive conditional is valid.[18] Then there is what they call restrictive obligation. A statement of restrictive obligation has the form "Given that p, S is obligated to do X," where this is understood as not being a conditional whose consequent concerns an unconditional obligation. As they again say, detachment with respect to such a statement is not valid.[19] They also note that ordinary English can be confusing, since both types of conditional obligation can be legitimately expressed by means of statements of the form "S ought to do X, if p."[20]

Now, it is quite correct to draw a distinction between hypothetical and restrictive obligation and to note that detachment is valid for the former but not the latter. This helps Jackson and Pargetter, though, only if it can be shown that my obligation to decline the task, if I won't provide comments on time, is a hypothetical rather than a restrictive one. This, of course, is what Jackson and Pargetter contend, but their rationale is problematic, for three reasons.

Consider, first, the voting case. There is an inconsistency in Jackson's and Pargetter's position with respect to this case. They declare the obligation that I ought to vote for Dick, if I don't vote for Tom, to be a restrictive one. This seems absolutely right. But then they say that, since detachment for such a conditional obligation is not valid, it is not the case that I ought to vote for Dick. Yet we have seen that,

on their version of actualism, I ought to vote for Dick (since my doing so would be better than my not doing so).

Second, their position overlooks the apparent isomorphism between the voting case, where it is granted by all concerned that restrictive obligation is at issue, and the commenting case. The mark of restrictive obligation, as Jackson and Pargetter rightly note, is that it concerns obligation relative to a restricted set of options. I ought to vote for Tom, given all my options; but, with the option of voting for Tom discounted, I ought to vote for Dick. Similarly, I ought to accept, given all my options; but, with the option of commenting discounted, I ought to decline. Or again: I ought to vote for Tom; if I don't vote for Tom, I ought at least to vote for Dick. Similarly: I ought to comment on the paper; if I won't comment on it, I ought at least (to have the courtesy) to decline the task. Jackson and Pargetter in fact admit that the propriety of using "at least" signals restrictive obligation.[21] Since it is proper to use it in the commenting case, again we find a tension in their position.

Third, their position is undermined by their own characterization of hypothetical obligation. They say that what follows "if" in cases of hypothetical obligation involves not a restriction of options but a re-ordering of them. This again seems quite right. They give this example: "I ought to vote for Dick if Tom is (were) a crook."[22] What this says is: if it is (were) true that Tom is a crook, then my options should be ranked not as previously but in some such way as follows:

    1: —T & D & —H
    2: —T & —D & H
    3: —T & —D & —H
    4: T & —D & —H.

But notice how different this is from the commenting case. By saying that I ought to decline the task if I won't comment on the paper, I am emphatically not saying that, since it is true that I won't comment, the ranking of my options is to be revised in some such way as follows:

    1: —A & —C
    2: A & —C
    3: A & C.

I think, then, that we simply must accept that the obligation to decline the task, if I won't comment, is a restrictive obligation. Since unrestricted detachment is acknowledged by all to be invalid for such obligation, we cannot argue that it is such detachment that warrants the claim that the actualist wants to make, namely, that I ought to decline the task.

Having just asserted that detachment for restrictive obligation is not valid, I now want to complicate matters by declaring that it is, in a sense, valid nonetheless. In this way, I shall seek to accommodate much, if not all, of what makes actualism appealing to many, at the same time as avoiding its shortcomings. Return to the voting case. We may agree that: I ought to vote for Tom; I ought not to vote for Dick; but if I don't vote for Tom, I ought (at least) to vote for Dick. The inclusion of "(at least)" is important. It indicates that, if I fail to vote for Dick as well, I will have compounded my wrongdoing. That is, there seems to be a sense in which, if I don't vote for Tom, then I do have an unconditional obligation to vote for Dick, such that, if I fail to vote (even) for Dick, I will have done an extra wrong.

But how can it be that I ought not to vote for Dick and that I ought to vote for Dick? The proper answer, I believe, is roughly this.[23] We should distinguish levels of obligation. My primary obligation is not to vote for Dick (but Tom). If I fail in this primary obligation, though, a subsidiary (here, second-level) unconditional obligation arises (by way of detachment from a conditional obligation), namely, to vote for Dick. There is no conflict here, though, of the sort ruled out by (T5) above, for the obligations are on different levels. Indeed, the story could be continued: if I fail to vote for either Tom or Dick, I ought (at least) to vote for Harry; and so, if I do not vote (even) for Harry, I will have committed a third-level wrong. It is most important to note that this way of accounting for the wrong done in not voting for Dick does not sanction wrongdoing in the way that actualism does; for my primary obligation is to accept the task of providing comments on Jones's paper. The fact that, under the circum-stances, accepting the task would compound a wrongdoing does not alter the fact that I ought to accept the task. Here possibilism and actualism differ, to the credit, I believe, of the former.

Jackson and Pargetter claim that there is a "positive advantage" to actualism.[24] It can account for the fact that, intuitively, if I decline the task, I meet some but not all of my obligations. For, they say, I ought to decline the task (even though I ought also to accept-and-comment). But we can now see that the possibilist who accepts the possibility of subsidiary obligation can also say that, if I decline the task, I meet some but not all of my obligations. I meet the subsidiary obligation to decline it; I fail to meet the primary obligation to accept it. Thus, contrary to Jackson and Pargetter's claim, actualism enjoys no advantage here.

Part II.ii
The second objection is this. I might turn to you for advice as to what to do. Knowing my propensity for procrastination, you might well advise me to decline the task of providing comments on Jones's paper. Are you advising me to act immorally? Are you hoping for this? Clearly not. Your advice is the right advice, and this implies that what you advise me to do is right.[25]

This objection cannot be accepted. It would prove too much, for it could just as easily be applied to the voting case. Knowing my propensity for overlooking the best qualified candidate, you might well advise me to vote for Dick (rather than Harry). This advice would be just as appropriate here as your advice to me to decline the task of providing comments on Jones's paper. But, we all agree, it is not the case that I ought to vote for Dick.

So where has the objection gone wrong? Presumably in its implicit reliance on these two theses:

> (T6) If S ought to advise S' to do X, then S' ought to do X.
> (T7) If S advises S' to do X, and S' ought not to do X, then S advises S' to act immorally.

With respect to the first thesis, all that needs to be said is this. Suppose that I will follow your advice as to whether I should accept or decline, but that nothing you could do would persuade me to provide the comments on time. In such a case, the best that you can do does not match up with the best that *I* can do. Under the circumstances, you can, by means of your advice, see to it either that I accept but do not provide comments or that I decline (and do not provide comments). We may assume that the latter course of action is better. Thus possibilism implies in this case that you ought to advise me to decline, while still implying that I ought to accept.[26] As for the second thesis: we should note that there is a sense in which it is true and a sense in which it is false. The sense in which it is true is this: if S advises S' to

do X, and S' ought not to do X, then there is some action which S advises S' to perform, and that action is wrong. The sense in which it is false is this: if S advises S' to do X, and S' ought not to do X, then S advises S' to perform a wrong action — where this is to be understood as implying that the wrongness of the action forms part either of the content of the advice or of the objective of the advisor. After all, it is not your objective (or hope), when advising me to decline, that I should act immorally; that is something which you take as given, that is, as beyond your remedy (though not mine). What you are aiming at is the minimization of immorality on my part, consistent with the restrictions on your influence under which you are operating.

Part II.iii
The third objection is this. Suppose that Jones could get comments on Wednesday from someone else if, but only if, I decline the task; it would be best if I commented, worst if no one did. It is clear that I ought not to delay his getting comments; but, by hypothesis, this is just what I would be doing if I accepted the task, for my so accepting would be, under the circumstances, my delaying his getting comments. Hence I ought not to accept.[27]

This objection rests on two dubious assumptions. The first is that my accepting the task would be my delaying the comments. The second is this:

   (T8) If S ought to do X, and doing X is doing Y, then S ought to do Y.

These assumptions are jointly unacceptable. Either my accepting the task is not my delaying the comments, but is related in some less intimate way (expressed by "by" in "I delay the comments by accepting the task");[28] or (T8) is false since, even if my accepting the task is my delaying the comments, this identity holds only contingently, and I ought not to allow it to hold. That is, the possibilist, working under the first assumption, would accept

   (T9) If S ought to do X, and he ought to allow his doing X also to be his doing Y, then S ought to do Y,

but there would be no reason for him to accept, indeed every reason to reject, (T8).

Part III
The actualist's criticisms of possibilism can thus be met;[29] and actualism violates several plausible theses concerning obligation which possibilism does not. Hence I favor possibilism. But what, more exactly, does possibilism say about where it is that I go wrong when I do go wrong?

Part III.i
It might at first seem that this question should be answered as follows. When a person is obligated to perform an action, that action is "scheduled" to be performed at a certain time T;[30] he goes wrong at T if he fails to perform the action as scheduled. More pedantically, the answer is this:

   (A1) S goes wrong at T in not doing X if and only if
       (i) S ought to do X at T, and
       (ii) S does not do X at T.

While this answer is adequate for all cases of what we may call immediate obligation (I ought now to accept the task right now, but I don't; so I go wrong right now. I ought now to vote for Tom right now, but I don't; so I go wrong right now), it is not adequate for all cases of what we may call remote obligation. Consider, for example, the following case of "self-imposed impossibility."[31] On Monday I borrow some money from Smith, promising to repay it on Wednesday; on Tuesday I gamble the money away, rendering myself unable to repay it on Wednesday. If "ought" implies "can," then it is not the case that I ought on Wednesday to repay the money. If obligation could not be remote, it would then never be the case that I ought to repay the money on Wednesday. This seems a preposterous result. To avoid it, one must either deny that "ought" implies "can" or admit the possibility of remote obligation. Thus, in the present case, we can say that I ought on Monday (in part because I can on Monday) repay the money on Wednesday. The "ought" remains in effect up until that time on Tuesday when I render myself unable to repay the money on Wednesday; at that time, a new set of obligations comes into effect (among which are, presumably, those of apologizing to Smith on Wednesday and of making amends to him as soon after Wednesday as I can).

We may assume that this case involves the following ranking of options:

    1: B(orrow) & —G(amble) & R(epay)
    2:—B & —G & —R
    3: B & —G 8c —R
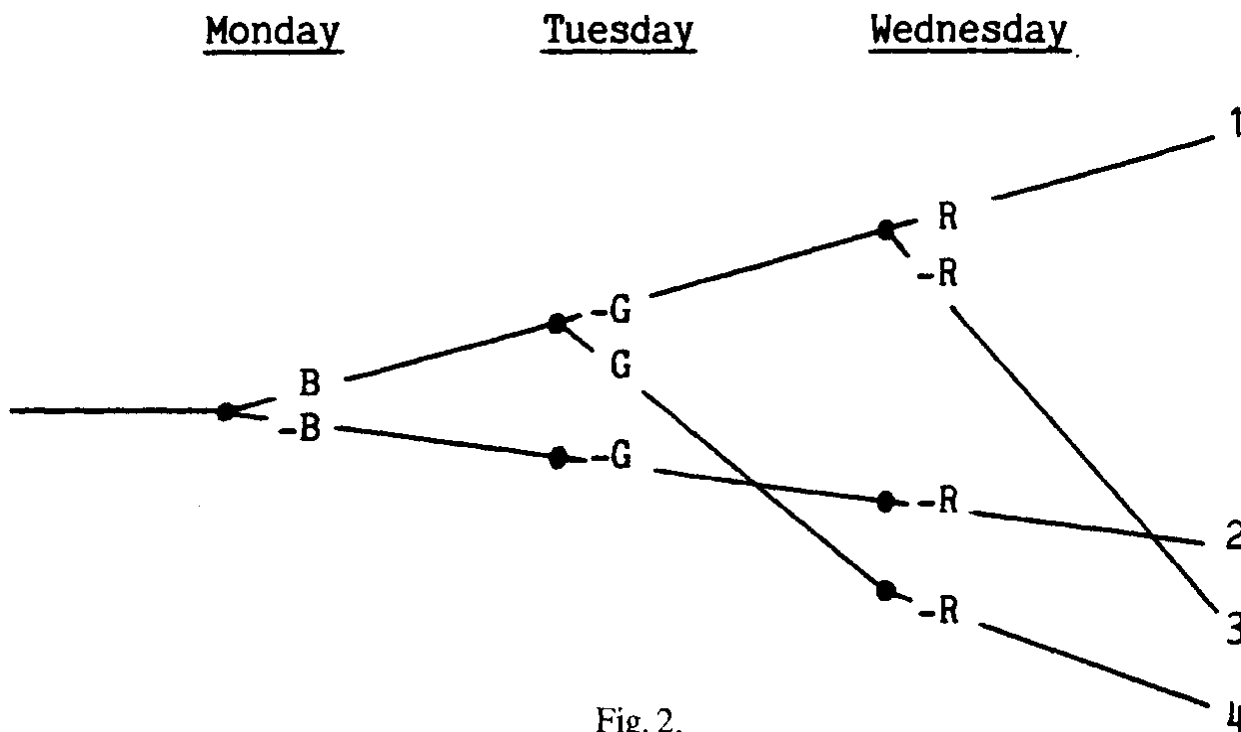   4: B & G & —R.

The picture is this:

Fig. 2.

Notice again that I ought on Monday to repay on Wednesday, but on Tuesday I gamble the money away

so that it is not true that I ought on Wednesday to repay on Wednesday. According to (A1), the wrong I do in not repaying is done on Wednesday. But how can it be that I do wrong in not repaying at a time when I am not obligated to repay? This seems incoherent.

Part III.ii
There are a number of ways to avoid such incoherence. One way is to deny the possibility of remote obligation. It might be said that what I do wrong is to gamble, and this is something I ought right then not to do. But this response ignores the fact that, while we may admit that the gambling is wrong, it is wrong because of the wrongness of not repaying. (That is, we may assume that the gambling would not have been wrong if there had been no obligation to repay.) Clearly, on Thursday, we want to say that I ought to have repaid (it was wrong of me not to repay) the money on Wednesday. If this is so, either remote obligation must be admitted as possible or, as noted before, the thesis that "ought" implies "can" must be jettisoned. Again, the former alternative seems preferable to me.

Since it was the gambling that rendered the repayment impossible, and since it has been admitted that the gambling was wrong because of this, it might seem that a better answer to our question is this: a wrong is done either at the time at which an obligatory action is "scheduled" to be performed or at the time at which this action is rendered impossible. That is:

> (A2) S goes wrong at T in not doing X at T' if and only if
> (i) S ought at T to do X at T',
> (ii) S does not do X at T' , and
> (iii) it is not the case that S can after T do X at T'.

But this won't do either. Consider this case.[32] On Monday I promise to help you move furniture on Wednesday; on Tuesday I kidnap a child, rendering it obligatory for me to return the child on Wednesday; while I can still help you move furniture on Wednesday, I cannot both do that and return the child. Let us assume that, due to my promise, I ought on Monday to help you move furniture on Wednesday; let us also assume that I return the child on Wednesday. Where did I go wrong in not helping you move furniture on Wednesday?

The situation is this (where, for simplicity's sake, it is assumed that I have already promised to help you):

> 1: —K(idnap) & H(elp) & —R(eturn)
> 2: —K & —H & —R
> 3: K & —H & R
> 4: K & H & —R
> 5: K & —H & —R.

The picture is this:
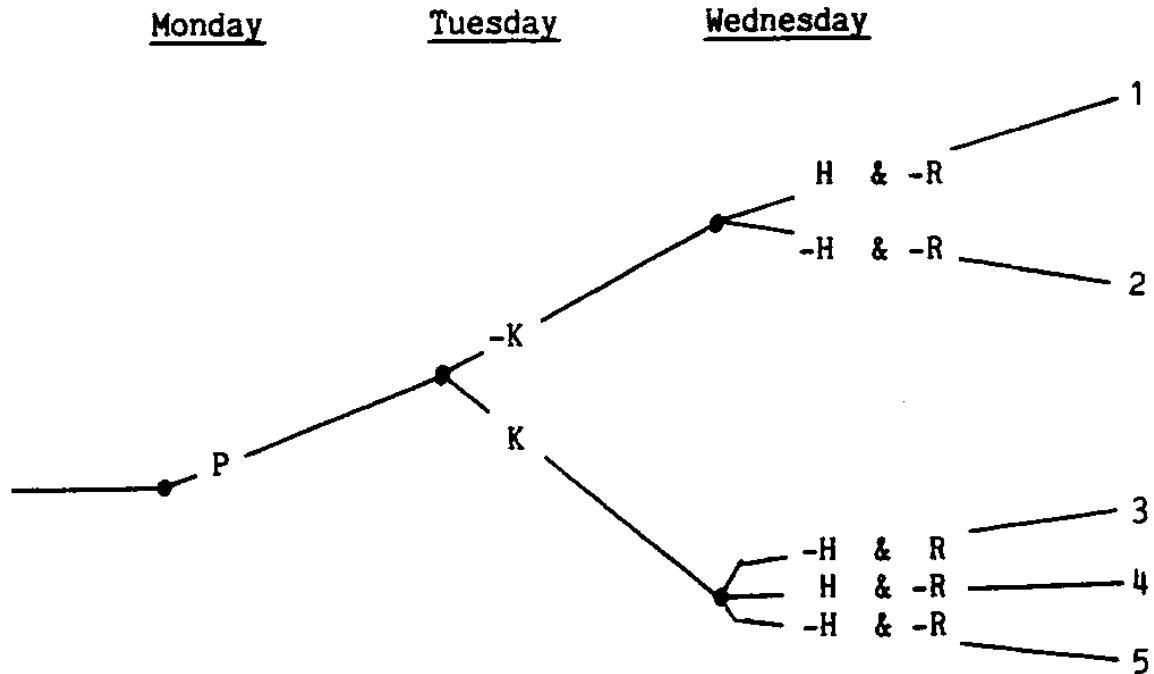
Monday   **Tuesday**   **Wednesday**

Fig. 3.

Clearly, on this picture, I ought on Monday to help you move furniture on Wednesday. By taking Track 3, I do wrong in not helping you. But I do not do wrong on Wednesday in not helping you, for Track 3 is preferable to Track 4. However, at no time prior to Wednesday does it become impossible for me to help you then. Therefore the present suggestion, (A2), concerning where I go wrong with respect to not helping you is unacceptable.

Part III.iii
The picture helps us answer the question properly, though. While I can still help you move furniture after I have kidnapped the child, if I do not do so I will violate an obligation which applied to me up until the time that I kidnapped the child, that is, up until that time on Tuesday where I veered from Track 1 on to Track 3. It therefore seems plausible to say the following:

> (A3) S goes wrong at T in not doing X at T' if and only if
>   (i) S ought at T to do X at T' ,
>   (ii) S does not do X at T' , and
>   (iii) it is not the case that S ought after T to do X at T'.

The suggestion here is this. To find out where an agent goes wrong with respect to not performing an action, retrace his steps back down the track which he in fact took until you reach a juncture, a point of choice, where it is true that he ought then to perform the action — where, that is, there is a top-level track leading out from that point and on which the action in question appears as performed. In Figure 3, we need only go back along Track 3 to the point of choice on Tuesday to find a point where it is true that I ought at that time to help you move furniture on Wednesday. Hence that is where I went wrong with respect to not helping you.

But, in fact, this won't quite do either. To see why, consider this case. I have promised to meet you at noon, and I ought to do so. But, in going to the meeting, I cause a minor accident at 11:30, and this requires me to make minor reparation (which I cannot do if I keep the appointment). In going to make reparation, however, I cause a major accident at 11:45, and this requires me to make major reparation (which I cannot do if either I keep our appointment or make reparation for the first accident). What I in fact do is go on to make reparation for the first accident. The situation is this:

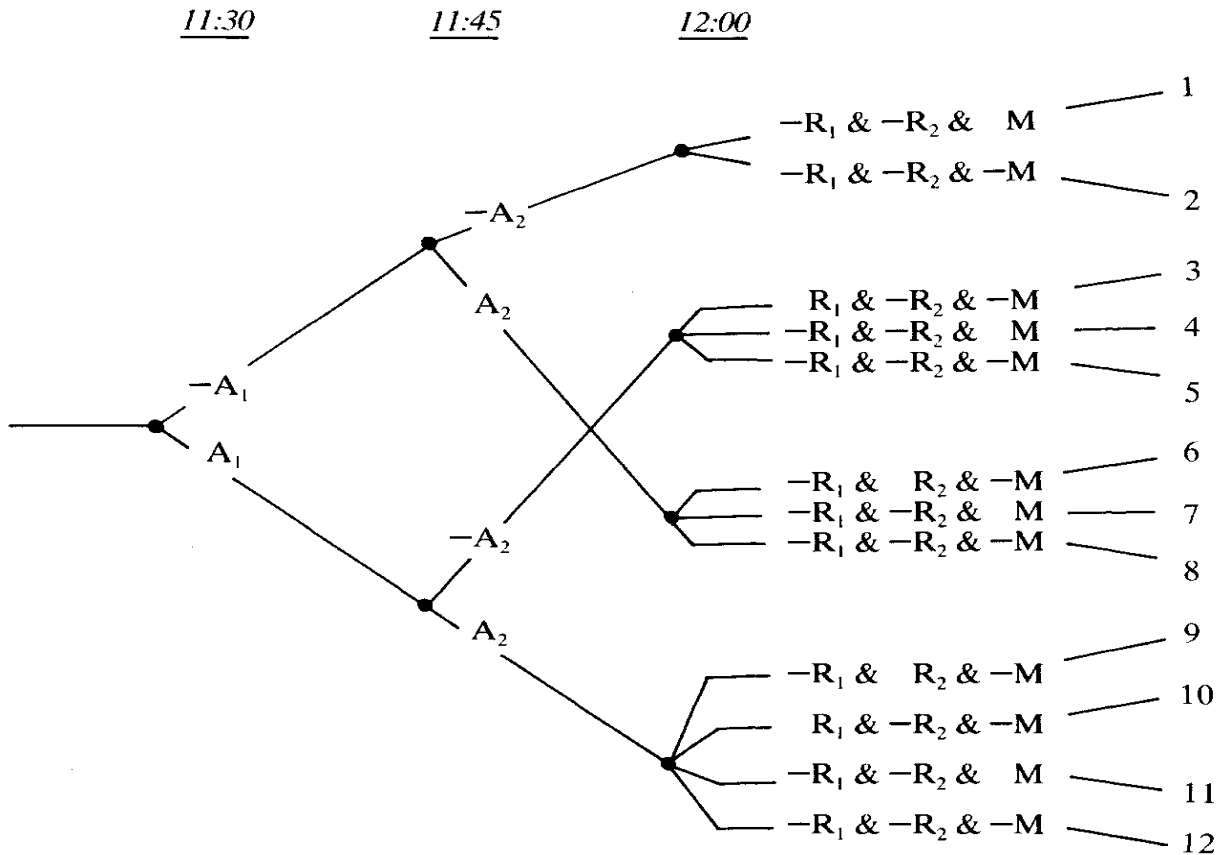|	
1:	$-A(ccident)_1$ & $-R(eparation)_1$ & $-A(ccident)_2$ & $-R(eparation)_2$ & $M(eet)$
2:	$-A_1$ & $-R_1$ & $-A_2$ & $-R_2$ & $-M$
3:	A_1 & R_1 & $-A_2$ & $-R_2$ & $-M$
4:	A_1 & $-R_1$ & $-A_2$ & $-R_2$ & M
5:	A_1 & $-R_1$ & $-A_2$ & $-R_2$ & $-M$
6:	$-A_1$ & $-R_1$ & A_2 & R_2 & $-M$
7:	$-A_1$ & $-R_1$ & A_2 & $-R_2$ & M
8:	$-A_1$ & $-R_1$ & A_2 & $-R_2$ & $-M$
9:	A_1 & $-R_1$ & A_2 & R_2 & $-M$
10:	A_1 & $-R_1$ & A_2 & $-R_2$ & M
11:	A_1 & $-R_1$ & A_2 & $-R_2$ & M
12:	A_1 & $-R_1$ & A_2 & $-R_2$ & $-M$.

And:

Fig. 4.

The track that I in fact take is Track 10. On this track, I perform R_1. In doing so, I both satisfy and fail to satisfy some obligations. I satisfy the obligation that arises out of my performing A_1, but I fail to satisfy the obligation not to perform A_1 (which would have brought with it not performing R_1), and I fail to satisfy the obligation to perform R_2 instead of R_1 (insofar as I wrongly performed A_2). The question presently at issue is this: where did I go wrong with respect to performing R_1? (A3) implies that I went wrong at 12:00, for I ought right at 12:00 not to perform R_1, I do perform it, and it is not the case that I ought after 12:00 not to perform it. Now, I do not wish to deny that I do go wrong at 12:00 in this regard. After all, I ought at 12:00 to take Track 9, not Track 10. Nonetheless, this is not the whole story with respect to where I go wrong in performing R_1 at 12:00. For, if we were to retrace my steps back down Track 10, we'd find a point of choice and a top-level track leading out of it on which my not performing R_1 appears. We find this, not at 11:45 (for the top-level track leading out from that point of choice is Track 3, and on this track I do perform R_1), but at 11:30. And so I want to suggest that I go wrong at 11:30 also in performing R_1 at 12:00. Or, in other words, I violate two obligations in performing R_1 at 12:00: my obligation at 11:30 not to perform it (an obligation which ceased at that point) and my new obligation (beginning at 11:45) not to perform it. Of course, in performing R_1 at 12:00 I also satisfy the obligation that I have, from just after 11:30 until 11:45, to perform it then.

In light of this, I suggest that the correct answer to our question is this:

> (A4) S goes wrong at T in not doing X at T' if and only if
>   (i) S ought at T to do X at T' ,
>   (ii) S does not do X at T' , and
>   (iii) it is not the case that S ought just after T to do X at T'.

My use of "just after" is intended to capture the point that, if there is a time after T at which S ought to do X at r, then there is an intervening time at which it is not the case that he ought to do X at T'.

It is a striking feature of the present proposal that whether and how an agent does go wrong at a certain time can hinge on what the agent will do at a later time. This is a feature which possibilism, as elaborated here, shares with actualism.[33] This feature may at first be somewhat disconcerting, but I submit that it is perfectly acceptable.

Part III.iv
Finally, we should consider this question. Must every failure to fulfill an obligation constitute wrongdoing? Suppose that, instead of my gambling my money away, it had been stolen from me. Would my failure to repay you on Wednesday then have been wrong? Or what if you had waived the debt? Would it still have been wrong for me not to repay you? Or suppose that the accident I had en route to our appointment was not my doing. Would my missing the appointment in order to help the victim of the accident then have been wrong? Or suppose that I did not kidnap the child but nonetheless found that I had the opportunity to return him to his parents. Would my failure to help you move furniture then have been wrong?

The proper answer here, from the perspective of possibilism, is this. If the events in question were not avoidable by me, then my obligations were not as they initially appeared to be. For example, if I could not have avoided losing the money, then I never was obligated to repay it on Wednesday. If I could not have avoided your waiving the debt (even if you could have avoided this), then, again, I never was

obligated to repay you. If I could not have avoided the accident, then I never was obligated to meet with you. If I could not have avoided the opportunity to return the child, then I never was obligated to help you move furniture. All this is of course quite consistent with its being reasonable for me (or anyone else) initially to believe that I was obligated in these ways.

On the other hand, if (as is far more likely) the events in question were avoidable by me, then either my obligations were not as they initially appeared or I was obligated to avoid the events. Either I was not obligated to repay you, or I was obligated to avoid losing the money (or to avoid your waiving the debt); either I was not obligated to meet you, or I was obligated to avoid having an accident; either I was not obligated to help you, or I was obligated to avoid the opportunity to return the child. In this way, we may stick to the thesis that every failure to fulfill an obligation does indeed constitute wrongdoing. But what if, for example, I could have avoided the accident (by taking Elm Street rather than Oak Street, say), but there was no reason for me to believe this (no reason to believe that an accident was in the offing at all), I took all reasonable precautions, I was not at fault? Did I still do wrong in not avoiding it? Yes, if I was obligated to meet you. I went wrong in turning down Oak Street. This is perfectly compatible with my not being at fault in so doing.[34] This paper is concerned with the wrongness of actions, not the faultiness of agents. It is a fundamental error to think that an agent's doing wrong implies that the agent is at fault. 'The link between wrongdoing and faultiness is far more indirect than that.[35, 36]

NOTES

1. These formulations of actualism and possibilism are rough. There need not be just one way that things would or could turn out if S did X. For a more precise rendition of possibilism, see [201, p. 66, which was itself inspired by pl. Since writing [3], Feldman has provided a very elaborate account of possibilism in [2].
2. See [5] and [14].
3. See [2], [6], [7], [8], and [15].
4. See [9].
5. For similar use of such a device, see [16], p. 171. See also [6], p. 194.
6. [9], p. 247.
7. That (T2) implies (T1) can been seen by substituting "X and Y' for "X" in (T2). Thus: if S ought to do both X and Y, and it is logically impossible (which it is) that he do X and Y without doing Y, then S ought to do Y. Since "X" and "Y' are inter-substitutable in the antecedent of this resultant formula, we may also infer that S ought to do X.
8. The proper construal involves certain time assignments plus an extra clause about ability. It is this (where T is no later than either T' or T"): if S ought at T to do X at T', and it is logically impossible that S do X at T' without doing Y at T", and S can at T not do Y at T", then S ought at T to do Y at r. (Corresponding restrictions can be placed on (T1).) It may seem that (T2), even when so understood, is still vulnerable to the following sort of case (taken in modified form from [1], pp. 214-215): Smith ought at T_1 to help at T_2 a man whom he will mug at T_3 (call this complex action X); his doing X entails his mugging the man at T_3 (call this Y); he can at T_1 refrain from doing Y at T_3; yet it is not the case that he ought at T_1 to do Y at T_3. But the claim that his doing X entails his doing Y and the claim that he can refrain from doing Y are in fact jointly unacceptable. Certainly there are possible worlds in which it is true that Smith helps at T_2 the man whom he mugs at T_3 in the actual world and in which it is false that Smith mugs the man at T_3. On this understanding of what Y is, then, we have no counterexample to (T2); for Smith's doing X does not entail his doing Y. On the other hand, it is true that there are no possible worlds

in which it is true that Smith helps at T_2 the man whom he mugs at T_3 in the actual world and in which it is false that Smith mugs the man at T_3 in the actual world. But it is also true that, given that Smith does mug the man at T_3 in the actual world, there is no possible world in which it is false that he mugs the man at T_3 in the actual world. (See [121, p. 55.) This implies, though, that, on this understanding of what Y is, Smith cannot at T_1 refrain from doing Y at T_3; and so, once again, we have no counterexample to (T2).

9. (T3) — which can be understood along the lines of (T2)'s construal, as given in the last note — implies (T2), because whatever is logically impossible is personally impossible.
10. See [191, section 11.
11. Compare [9], p. 249.
12. [9], p. 236.
13. [9], p. 247.
14. [10], pp. 77-78.
15. While I have not shown it, the reader may easily confirm that possibilism in fact implies that each of (T1) through (T5) is true.
16. [9], p. 238.
17. [10], pp. 77-78. They also hint at it in note 20 on p. 253 of [9]. Compare [20], p. 67.
18. [10], p. 77. Compare [2], p. 82, on "subjunctively conditioned absolute moral obligation."
19. [10], p. 78. Compare [2], p. 90, on "conditional moral obligation."
20. [10], p.80.
21. [9], p. 253, n. 20.
22. [10], p. 79.
23. I give a fuller answer in [20], an answer itself inspired by [11] and anticipated in [6]. This sort of answer is mentioned but presumably rejected by Jackson and Pargetter in [9], p. 244, n. 11.
24. [9], p. 241.
25. [9], p. 237.
26. Compare [2], pp. 55-57.
27. [9], p. 239.
28. Compare [4], Chs. 1 and 2, and [17], Ch. 7.
29. Jackson and Pargetter in fact offer a fourth objection, claiming that possibilism is arbitrary ([9], p. 237). I do not understand the objection, and so I shall not discuss it here.
30. I shall ignore complications concerning the "scheduling" of an action by, or within, rather than at a certain time.
31. This case is taken in modified form from [19], p. 199.
32. This case is taken in modified form from [19], pp. 201-202.
33. It is an instance of what is called covert relationality in [13], pp. 121-125. (Contrast, to some extent, the final sentence of [6].)
34. Compare [15], pp. 178-179.
35. I examine this link in [18], at least where faultiness is understood in terms of moral culpability.
36. This paper was presented in January 1989 at a meeting in Chapel Hill of the Research Triangle Ethics Circle and has benefited from the comments of those present. I especially wish to thank David Copp and Thomas E. Hill, Jr., for their help. I am grateful also for the helpful remarks of an anonymous referee.

BIBLIOGRAPHY
[1] Castatieda, Hector-Neri. *Thinking and Doing*. Dordrecht: D. Reidel, 1975.

[2] Feldman, Fred. *Doing the Best We Can*. Dordrecht: D. Reidel, 1986.

[3] Feldman, Fred. "Oughts — Absolute, Conditioned, Conditional." *Philosophia*, 12 (1983), 257-272.

[4] Goldman, Alvin I. *A Theory of Human Action*. Princeton: Princeton University Press, 1976.

[5] Goldman, Holly S. "Dated Rightness and Moral Imperfection." *The Philosophical Review*, 85 (1976), 449 —487.

[6] Goldman, Holly S. "Doing the Best One Can." In A. I. Goldman and J. Kim, eds., *Values and Morals* (Dordrecht: D. Reidel, 1978), 185-214.

[7] Greenspan, Patricia S. "Oughts and Determinism: A Response to Goldman." *The Philosophical Review*, 87 (1978), 77-83.

[8] Humberstone, I. L. "The Background of Circumstances." *Pacific Philosophical Quarterly*, 64 (1983), 19-34.

[9] Jackson, Frank and Pargetter, Robert. "Oughts, Options, and Actualism." *The Philosophical Review*, 95 (1986), 233-255.

[10] Jackson, Frank and Pargetter, Robert. "Two Puzzles about Conditional Obligation." *Philosophical Papers*, 16 (1987), 75-83.

[11] McKinsey, Michael. "Levels of Obligation." *Philosophical Studies*, 35 (1975), 385-395.

[12] Plantinga, Alvin. *The Nature of Necessity*. Oxford: Clarendon Press, 1974.

[13] Slote, Michael. *Common-sense Morality and Consequentialism*. London: Rout-ledge and Kegan Paul, 1985.

[14] Sobel, Howard. "Utilitarianism and Past and Future Mistakes." *Nous*, 10 (1976), 195-219.

[15] Thomason, Richmond H. "Deontic Logic and the Role of Freedom in Moral Deliberation." In R. Hilpinen, ed., *New Studies in Deontic Logic* (Dordrecht: D. Reidel, 1981), 177 —186.

[16] Thomason, Richmond H. "Deontic Logic as Founded on Tense Logic." In R. Hilpinen, ed., *New Studies in Deontic Logic* (Dordrecht: D. Reidel, 1981), 165— 176.

[17] Zimmerman, Michael J. *An Essay on Human Action*. New York: Peter Lang, 1984.

[18] Zimmerman, Michael J. *An Essay on Moral Responsibility*. Totowa: Rowman and Littlefield, 1988.

[19] Zimmerman, Michael J. "Remote Obligation." *American Philosophical Quarterly*, 24 (1987), 199-205.

[20] Zimmerman, Michael J. "Subsidiary Obligation." *Philosophical Studies*, 50 (1986), 65-75.