

Initiating Coordination

Paul Weirich^{†‡}

How do rational agents coordinate in a single-stage, noncooperative game? Common knowledge of the payoff matrix and of each player's utility maximization among his strategies does not suffice. This paper argues that utility maximization among intentions and then acts generates coordination yielding a payoff-dominant Nash equilibrium.

1. Introduction. Committee members meet to conduct business. Drivers stay on the right side of the road. Workers divide the task of painting a house. These are all examples of coordination. What principles of rationality lead individuals to do their parts in a method of coordinating? Providing an answer presents a challenge for game theorists.

2. The Hi-Lo Paradox. This paper treats only coordination in single-stage noncooperative games. In such games the players do not causally interact, so coordination requires an epistemic foundation. Some theorists hold that classical principles of rationality cannot provide that foundation.

Suppose that agents are *hyperrational* in Howard Sobel's (1994) sense. They are moved only by utility maximization. Habit, convention, precedence, salience, focal points, and so on have no effect unless they affect the utilities of options. These common coordination devices are ineffective among hyperrational agents. Such agents respond only to another's participation in an equilibrium. None moves toward an equilibrium unless the others do. None gets the ball rolling, as Robert Sugden (2000) and Margaret Gilbert (2001) observe.

Andrew Colman (2003) illustrates the problem using the two-person, pure-strategy game Hi-Lo, which the following payoff matrix (Table 1) depicts. There are two forms of coordination, one offering a high payoff to each player, and another offering a low payoff to each player. A player

[†]To contact the author, please write to: Philosophy Department, University of Missouri, Columbia, MO 65211; e-mail: WeirichP@missouri.edu.

[‡]I thank the audience at my paper's presentation at the 2006 PSA meeting for many insightful points.

TABLE 1. THE TWO-PERSON,
PURE-STRATEGY HI-LO GAME.

	High	Low
High	2, 2	0, 0
Low	0, 0	1, 1

has no grounds for assigning probabilities to his partner's strategies. Classical decision principles give the players no reason to do their parts in the superior equilibrium (High, High). Indeed, utility maximization grounds neither coordination equilibrium. It fails even granting common knowledge of the payoff matrix and of the players' utility maximization. Individual rationality appears to not generate coordination in Hi-Lo. That is the Hi-Lo paradox.¹

This paper resolves the Hi-Lo paradox. It shows that individual rationality generates coordination in the game Hi-Lo, at least if agents are ideal and in ideal circumstances. Unlike the psychological game theory Andrew Colman advances, it normatively supports coordination in ideal cases. The idealizations reveal basic principles of strategic rationality forming the structure of a general theory that eventually rescinds the idealizations.

In contrast with some approaches to coordination, my approach does not rely on repetition of games or the bounded rationality of agents. It does not offer evolutionary explanations of coordination as, for example, Brian Skyrms (2004) does. My objective is to explain the rationality of realizing the superior equilibrium in coordination problems such as the game Hi-Lo. This requires an account of the strategic reasoning of agents. By way of analogy, consider equilibrium for a ball rolling in a basin with several low spots. It is easy to show that the ball will settle in a low spot. An explanation of its settling in a particular low spot is more complex, however. It provides an account of the ball's travels before coming to rest. Similarly, an explanation of an equilibrium's realization does not show just that an equilibrium is realized but also provides an account of deliberations leading to the equilibrium's realization. An explanation accounts for the agents' strategic reasoning and does not just characterize its outcome.

This paper explains how strategic reasoning in coordination problems leads agents to a Nash equilibrium, in particular, a payoff-dominant Nash equilibrium. Of course, explanation must start somewhere, and I make some assumptions about the agents' knowledge without explaining how the agents acquire that knowledge. I show how strategic reasoning gen-

1. Donald Regan (1980, 18) presents a similar problem of coordination for two agents committed to the utilitarian principle of morality.

erates coordination given that agents have some initial knowledge about themselves and their coordination problem.

3. Commitment. A common view suggests that coordination may arise from commitment to coordination and to superior forms of coordination when there are multiple forms of coordination. This section examines commitment as a method of achieving coordination. Subsequent sections evaluate specific applications of the method.

In noncooperative games commitment influences others only if they discover the commitment despite the absence of communication. The discovery may arise from the players' initial knowledge of the game and of the players. People readily commit to coordination in games such as Hi-Lo and others know that they do. Knowledge of the payoff matrix of Hi-Lo is often enough to lead each player to High. Does coordination arising this way involve any irrationality?

Consider first, hyperrational agents. Unlike humans, hyperrational agents do not perform an act unless it maximizes utility. In Hi-Lo the act High does not maximize utility unless one's counterpart performs it too. Hyperrational agents are stuck. Neither player can use commitment to initiate coordination. A hyperrational agent does not have recourse to a pill that induces High, for instance. Although it may be rational to take such a pill, it induces behavior not guided by utility maximization. Taking such a pill is inconsistent with hyperrationality by definition. Any commitment to perform High raises the probability of High. By definition, a hyperrational agent increases the probability of High only if High maximizes utility. So he does not form the commitment. The commitment to perform High has good consequences, but hyperrationality blocks the commitment.

May rational agents, in contrast with hyperrational agents, use commitment to initiate coordination? May any of the commitment devices that humans use, for instance, be applied without irrationality at any step? Commitments come in various types. Perhaps some type of commitment resolves the Hi-Lo paradox without any irrationality.

Without communication, agents may know the agreement they would reach if they were to communicate. That agreement is a focal point of the sort used by Thomas Schelling (1960) to explain coordination. Each may decide to do his part in the agreement. A disposition to decide this way may be rational to form and to have. The act it issues is not irrational even if performed in ignorance of its consequences. It is as good as alternative acts, given ignorance of the acts of others. Moreover, agents with common knowledge of the focal point may expect each other to do their parts in that strategy profile, and then doing their parts maximizes utility.

The appeal to focal points just sketched suggests a general method of using commitment to justify coordination. Suppose that agents begin the game Hi-Lo already equipped with a disposition to participate in its payoff-dominant Nash equilibrium. Knowledge about agents may yield evidence about the strategy profile they will realize. Each agent is committed to the payoff-dominant equilibrium. The other agent knows this. The commitment is rational to acquire because each agent knows that the other responds to the commitment by doing her part in the payoff-dominant equilibrium. So each agent knows that he will not act alone. Commitment gets the ball rolling toward coordination.

Is an agent's commitment to High rational to have? It is a commitment to do his part in the payoff-dominant Nash equilibrium without regard for his partner's act. The commitment yields an irrational act if the agent knows his partner will perform Low. The agent may know that the adverse circumstances do not obtain, however. The commitment's disregard for reasons may not lead to an act that flies in the face of reason. So perhaps in favorable cases the commitment is rational to maintain as well as to form.

Is it rational for an agent to perform the act the commitment issues? The agent commits to his part in the payoff-dominant Nash equilibrium without independent knowledge of his partner's part. Given his commitment, he expects his partner to do her part. Doing his part in the payoff-dominant equilibrium then maximizes utility because it arises from a disposition his partner knows he has, and knowing that she knows is evidence that she will do her part in the payoff-dominant equilibrium.

This is the general method of coordination through commitment. It alleges that pre-game commitment resolves the Hi-Lo paradox for comprehensively rational agents given that they are informed about each other and their game. Making the method work requires specifying the type of commitment and verifying that it leads to coordination without any missteps.

To illustrate the benefit of commitments, consider contracts and promises. They may be utility maximizing even if they sometimes issue acts that do not maximize utility. They may be rational gambles. Commitment to perform an act without independent knowledge that the act is utility maximizing need not result in a violation of the principle of utility maximization. The commitment may be a rational gamble, and the act itself may maximize utility. In Hi-Lo, an agent's commitment to High may not result in a violation of the principle of utility maximization. Whether it does, depends on his information about his partner's response to his commitment. Furthermore, whether his performing High is rational depends on his information about his partner's act. A rational act does not require knowledge that it is utility maximizing. Given ignorance, ratio-

nality requires only that it maximize utility under a quantization of beliefs and desires. An agent does not know High's utility ranking without knowing his partner's act. His performing High may maximize utility under a quantization of his beliefs and desires. Given ignorance of his partner's act, his performing High may be rational.

Although noncooperative games are causally static, they have inferential dynamics. A player may have a commitment to a payoff-dominant Nash equilibrium. Knowledge of that commitment may trigger another player's participation. The players' knowledge of their participation may make it utility maximizing for each to participate. An agent's commitment may be self-enforcing in the sense that if others know about it he has no reason not to honor it. If others know what he will do and respond by following his lead, and he knows that they will, then he has a reason to honor his commitment. Honoring it achieves coordination.

A commitment's formation may be rational because of its good consequences rather than because of the good consequences of the act it issues. Consider a commitment to perform an act regardless of the act's consequences. The commitment may be rational because of its consequences despite its being defectively motivated by the act's consequences. The commitment may be fully rational to have despite the defect in motivation because circumstances may make the defect acceptable. Whether the commitment is fully rational depends on its specification and context. In favorable circumstances it may be fully rational.

The foregoing appeal to commitment is an incomplete sketch. Commitment is vague and metaphorical. A commitment is a disposition to choose a certain strategy, but there are many mechanisms for commitment. Whether a form of commitment is rationally grounded also depends on the details of rationality's principles. Section 4 considers whether team-reasoning uses an acceptable form of commitment. Sections 5 and 6 consider whether the commitment involved in an intention to perform an act rationally yields coordination. For each type of commitment yielding coordination, I investigate whether fully rational agents may form, maintain, and honor the commitment.

4. Team Reasoning. This section examines one way of fleshing out commitment to coordination. It considers commitment to a team. This type of commitment arises through identification with a group and involves a type of reasoning called team-reasoning. After describing the way team-reasoning grounds coordination, I ask whether each step of the process is fully rational.

Michael Bacharach (2006) uses group identification and team-reasoning to resolve the Hi-Lo paradox. He argues for realization of the payoff-dominant Nash equilibrium. His approach, like mine, is classical and

appeals to the players' strategic reasoning. The classical approach, although theoretically attractive, has seemed unworkable. Bacharach ingeniously shows how to overcome obstacles to it. He revises principles of reasoning so that group identification yields coordination without violating the principles. His appeal to group identification is carefully grounded in social psychology. However, his principles of reasoning are controversial because they reject individualism.

Bacharach begins by attending to the role of frames in decision-making. Different agents see the same decision problem differently, and the difference in their perspectives explains differences in their decisions. Framing also explains focal points and so elaborates Thomas Schelling's (1960) suggestions concerning coordination.² The way an agent frames a coordination problem affects her behavior in the problem. An agent may see herself as acting independently of others, or as acting as a member of a team. The latter perspective resolves the Hi-Lo paradox. If an agent does her part in the strategy profile that best advances team goals, she does her part in the payoff-dominant equilibrium.

Bacharach turns to social psychology for an account of the features of social interaction that trigger group identification. He argues that the game Hi-Lo has those features. Each player may be expected to identify with the pair of players. Then team-reasoning leads each player to High. Group identification may not resolve every coordination problem, but it yields the superior form of coordination in Hi-Lo.

Bacharach introduces restricted and circumspect team-reasoning to take account of an agent's beliefs about the prospects of other agents' adopting team-reasoning. The hedged forms of team-reasoning respond sensibly to the danger that others will exploit team-reasoners. They describe assurances about others that sensible team-reasoning requires.

Bacharach holds that rationality permits, although it does not require, team-reasoning. Is that reasoning fully rational? Team-reasoning permits an agent to act contrary to her preferences when they are fully rational and balance all considerations. Such acts violate a basic principle of rationality. Bacharach holds that an agent adopting a team perspective does not see herself as acting against her preferences. She frames her decision just one way, namely, as advancing the team's goals. However an agent may be simultaneously aware of both her own and collective preferences and may see that advancing collective preferences conflicts with her own preferences. In any case, even if an agent with a team perspective loses sight of her own preferences, the rationality of acting contrary to her own preferences is hard to defend.

2. Frederic Schick (2003) makes related points about framing.

Bacharach appeals to evolution to explain the origin of team-reasoning. He adopts Elliot Sober and David Wilson's (1998) argument for group selection and then claims that team-reasoning is evolution's mechanism for group selection among humans. Team-reasoning is the proximate cause of coordination. Bacharach notes that evolution is likely to yield efficient mechanisms of group selection and thinks that team-reasoning is efficient. However, a single form of reasoning, such as utility maximization, that relies on payoff transformations to achieve coordination is more efficient than two independent forms of reasoning, such as individual- and team-reasoning, together with a device for selecting a form of reasoning. Efficiency suggests that reasoning follows preferences, and that preference revisions assist collective objectives.

Bacharach claims that payoff transformations do not resolve the Hi-Lo paradox, but he has in mind restricted types of payoff transformation.³ Payoff transformations that make High a dominant strategy for each player resolve the Hi-Lo paradox. Certain types of payoff transformation reconcile team behavior with individual preferences. Commitment to a team may generate conciliatory transformations, and through them coordination, without abandoning individualistic standards of rationality.

When using payoff transformations and the commitments they generate to justify coordination, the crucial step is establishing the transformations' rationality. The classical tradition, which I am following, puts aside evaluation of basic preferences. Therefore, rather than launch an evaluation of payoff transformations, I turn to another account of commitment and coordination that is also individualistic but does not rely on substantive principles of rational preference formation.

5. Intentions. An intention to act is a type of commitment. In the game Hi-Lo, may an effective intention to do one's part in the superior equilibrium replace a noncognitively grounded disposition to do one's part in that equilibrium? An intention requires less preparation than other dispositions. An agent directly controls his intentions, whereas he only indirectly controls his other dispositions. This section and the next argue that fully rational agents may use intentions to generate coordination.

To start, I examine intentions. An intention is a short plan. So to elucidate intentions, I review some basic points about plans. Why does an agent adopt a plan? To adopt a plan is to decide now about future acts. Why not put off decisions about future acts until the times for them arrive? May it maximize utility among possible decisions at a time to adopt a plan rather than postpone decisions about future acts?

3. See Colman (2003) for a discussion of the types of payoff transformations Bacharach considers. See especially his comments on the "bloated" Hi-Lo game.

Cognitively limited agents have good reasons to adopt plans. Adopting a plan is a way of achieving an act. This way of realizing an act may reduce deliberation costs. Instead of deliberating at each moment, an agent just executes her plan to perform the act. Earlier deliberations settle her future act. Because of the costs of continuous deliberation, cognitively limited agents profit from using plans to realize acts.⁴

Imperfect agents such as humans also use intentions to improve the rationality of their acts. A person forms an intention to perform an act to raise the probability of performing the act. Deliberating to identify a rational act and then forming an intention to perform the act increases the probability of performing a rational act.

An ideal agent has no cognitive limits but still may gain from forming intentions. Forming an intention is a way of creating a reason to act, as Bruno Verbeek (forthcoming) observes. Intention formation therefore has at least two types of benefit. First, an intention may prevent vacillation, which is costly. For example, in the case of Buridan's ass, an intention tips the scales toward a pile of hay. It ends vacillation. Second, as the next section argues, forming an intention may initiate coordination. Even a perfect agent, who may act in a utility maximizing way without deliberation or decision, may profit from forming an intention to act because the intention facilitates coordination.

Forming an intention to perform an act is a type of commitment because knowingly failing to fulfill a current intention one can fulfill at will is incoherent. Having an intention to perform an act increases the subjective probability of the act. It gives oneself a reason to perform the act. A rational agent may form an intention to act because of the consequences of the intention rather than because of the consequences of the act. However, intentions require justification by their content. Whereas a disposition to act may be rational to have because of its effects alone, an intention to act is rational to have only if it has both good effects and cognitive support. That is why Greg Kavka's (1983) toxin puzzle is hard to solve. The intention to drink the toxin, despite its good consequences, is irrational unless there are good reasons to drink the toxin.

To show that a fully rational agent may form an intention to act, with the objective of giving himself a reason to act, consider an agent who is indifferent between his options. He forms an intention to perform a certain act. Forming the intention itself has good consequences. It ends the decision problem. The intention also meets cognitive standards. There is no reason to rescind it. Furthermore, there are reasons for the act intended. Given the intention, the act has good consequences. It fulfills an intention.

4. Michael Bratman (1999, Chapter 2) makes similar points about the role of planning for cognitively limited agents.

The act does not promise better consequences than its rivals prior to the intention, but after forming the intention the act promises better consequences than its rivals because it alone fulfills an intention. That extra reason tips the balance in favor of the act intended.

In a decision problem, options are possible decisions. A decision is the formation of an intention to perform an act. The decision is certain to be realized if adopted, but the act decided upon may not be realized despite the intention to perform it. An intention to act may not be carried out. Events may frustrate an intention, or an agent may abandon an intention.

According to a common view, which this section adopts, given full information a plan's execution is rational only if its execution is optimal. Given uncertainty a plan's execution is rational only if its execution maximizes utility. Similarly, given uncertainty a plan's adoption is rational only if its adoption maximizes utility. Its adoption's outcome typically but not necessarily includes the plan's execution. Because plans adopted may not be executed, the utility of a plan's adoption is not necessarily the same as the utility of its execution. A plan rational to adopt may be irrational to execute. Someone who resolutely executes a plan to retaliate if attacked acts irrationally, if he is attacked and retaliation brings no benefits, even though the plan's adoption was rational because of the prospect of deterrence.

Suppose that an agent has irrationally adopted a plan. Its irrational adoption is an unacceptable mistake. The execution of the plan is rational only if it is rational given correction of that mistake and so, typically, only if it is rational independently of the plan.

A plan rationally adopted is rationally retained if no new relevant consideration arises, such as a change in the beliefs or desires supporting the plan. It is irrationally retained if it is unresponsive to relevant changes in circumstances. Irrationally retaining a plan is also an unacceptable mistake.

Under what circumstances is a plan irrationally retained? Generally, persevering with a plan is the best way to reach the long-term objective that prompted the plan. Abandoning or changing a plan is usually inferior to persevering. For example, if one wants to go to Chicago and plans to take the train, then upon arrival at the train station it is generally better to persevere than to adopt a new plan to drive to Chicago. Only salient changes should trigger a review. The appropriate level and type of responsiveness is achieved, not by deliberation, but by habits governing spontaneous review of plans. One should cultivate habits of spontaneously reopening deliberations at appropriate moments. Reopening deliberations too readily is the mistake of distraction, or irresoluteness. Excessive resistance to reopening deliberations is the mistake of obstinacy, or inflexibility. The optimal habit reopens deliberations when the benefits outweigh

the costs. The benefits, like the benefits of gathering information, are increases in the expected value of maximizing utility. Calculating costs and benefits is costly. An optimal habit responds to them without calculation. Nonoptimal habits, easier to inculcate, respond to them less reliably. A rational agent develops the best habits his circumstances permit. A decision about execution of an act planned may be rational despite deliberation unjustified by benefits, or despite lack of deliberation justified by benefits. It is rational if the agent's habits of spontaneously reopening deliberation are rationally aimed at reopening deliberation if and only if justified by benefits. A fully rational agent makes no mistakes in retention and execution of an intention.

6. Intentions Yielding Coordination. This section shows that rational agents may solve coordination problems because they form intentions that all foresee. Moreover, comprehensively rational agents may solve coordination problems this way because the intentions yielding coordination are rational to form, maintain, and execute.

In a single-stage noncooperative game such as Hi-Lo, an intention prompts coordination not by causally influencing other agents but rather because other agents foresee it. Foreknowledge of the intention influences other agents. Others may know the agent has or will form the intention. He may announce his intention prior to the game, or others may know his character and aptness to form the intention, perhaps, because of his comprehensive rationality. This foreknowledge may causally influence the players' acts. In the game Hi-Lo all may know of at least one player that he is an instigator of coordination. Moreover, this may be common knowledge. The instigator then gives himself a reason to perform High by forming the intention to perform High.

An intention to do one's part in High may replace a noncognitively grounded disposition to do one's part in High. The intention requires less preparation than the disposition. Its origin relies on decision preparation less than the disposition does. Reliance on intention is an attractive alternative to reliance on other forms of decision preparation. Intentions furnish a type of commitment an agent easily controls.

Suppose that a player forms the intention to perform the strategy High. Full rationality requires the rationality of forming, maintaining, and executing the intention. It requires more than utility maximization among strategies in the game. It evaluates as well the player's intention to perform High. The intention requires grounding in reason and in that respect differs from a disposition induced by a pill. Without such grounding, the intention provides an insufficient bootstrapping reason for an act, as Michael Bratman (1987, 24–27) observes.

A player's intention to perform High has both good consequences and

cognitive support assuming that the other player foresees it. Given foreknowledge of the intention, the other player maximizes utility by performing High herself. Forming the intention to perform High is rational because it triggers coordination. The intention to perform High is cognitively supported because performing High maximizes utility given that one's partner performs High. Having the intention is cognitively supported by the act's benefits given the intention's effects.

Compare forming that intention with forming a belief for pragmatic reasons. Consider, say, acquiring the belief that one will win a contest. One acquires the belief so that one will obtain the performance benefits of confidence. To be rational, the belief must also be epistemically justified. So suppose also that one knows confidence in victory makes one perform well enough to win. Then although one forms the belief for pragmatic reasons, one's holding it is epistemically justified because having the belief raises its prospects of being true. It is a self-fulfilling prophecy.

For another analogy, compare forming the intention to perform High with a self-supporting belief. Suppose that I believe that I formed a belief today. Perhaps before I formed the belief I had not formed a belief today. After I form the belief, the belief supports itself. A rational person may form the belief without reason but may then retain the belief for epistemic reasons because the belief is self-supporting. The belief's rationality does not require supporting evidence prior to its formation because it provides its own support.

Forming the intention to perform High is utility maximizing given its consequences. Having the intention is also rational because of its consequences. The reasons for having the intention do not arise from the consequences of performing High but from the consequences of the intention itself. Furthermore, performing High is utility maximizing given the intention's formation and the response it elicits. Fully rational agents may form self-supporting beliefs and intentions because of their good consequences.

A hyperrational agent cannot use an intention to initiate coordination. Such an agent forms an intention to perform an act only if the act maximizes utility. An intention to perform High raises the probability of High. By definition, a hyperrational agent increases the probability of High only if the act maximizes utility. So he does not form the intention. The intention to perform High has good consequences, but his character prevents it. He does not form the intention to perform High because the act does not maximize utility. It does not maximize utility unless he forms the intention to perform High and thereby elicits High from his partner. He has a chicken-and-egg problem.

Fully rational ideal agents, even perfect agents, can generate reasons by forming intentions prior to games they enter. One creates a reason to

perform an act by intending to perform the act. Fulfillment of an intention is a reason to act. An agent's intention to perform High by itself does not yield a sufficient reason to perform High. However, if it gives his partner a reason to perform High herself, then it gives him a sufficient reason to perform High. For rational ideal agents informed about their game and each other, reasons to perform High are mutually reinforcing. Having the intention to perform High furnishes a new reason to perform High. Knowledge of the intention gives one's partner a reason to perform High. Her performing High gives one a reason to perform Hi, and so on. The intention increases the probability of Hi, and an increase is all it takes to get the ball rolling toward coordination.

Fully rational agents may generate reasons by forming intentions. They may form, have, and execute intentions to do their parts in optimal methods of coordination. The commitments their intentions constitute resolve coordination problems.

REFERENCES

- Bacharach, Michael (2006), *Beyond Individual Choice: Teams and Frames in Game Theory*. Princeton, NJ: Princeton University Press.
- Bratman, Michael (1987), *Intentions, Plans, and Practical Reason*. Cambridge, MA: Harvard University Press.
- (1999), *Faces of Intention*. Cambridge: Cambridge University Press.
- Colman, Andrew (2003), "Cooperation, Psychological Game Theory, and Limitations of Rationality in Social Interaction", *Behavioral and Brain Sciences* 26: 139–198.
- Gilbert, Margaret (2001), "Collective Preferences, Obligations, and Rational Choice", *Economics and Philosophy* 17: 109–119.
- Kavka, Gregory (1983), "The Toxin Puzzle", *Analysis* 43: 33–36.
- Regan, Donald (1980), *Utilitarianism and Co-operation*. Oxford: Clarendon.
- Schelling, Thomas (1960), *The Strategy of Conflict*. Cambridge, MA: Harvard University Press.
- Schick, Frederic (2003), *Ambiguity and Logic*. Cambridge: Cambridge University Press.
- Skyrms, Brian (2004), *The Stag Hunt and the Evolution of Social Structure*. Cambridge: Cambridge University Press.
- Sobel, J. Howard (1994), *Taking Chances: Essays on Rational Choice*. Cambridge: Cambridge University Press.
- Sober, Elliot, and David Wilson (1998), *Unto Others: The Evolution and Psychology of Unselfish Behavior*. Cambridge, MA: Harvard University Press.
- Sugden, Robert (2000), "Team Preferences", *Economics and Philosophy* 16: 175–204.
- Verbeek, Bruno (forthcoming), "Rational Self-Commitment", in Fabienne Peter and Hans Schmid (eds.), *Rationality and Commitment*. Oxford: Oxford University Press.