Labor union members play an OLG repeated game

Michihiro Kandori[a,1] and Shinya Obayashi[b]

[a]Faculty of Economics, The University of Tokyo, Tokyo 113-0033, Japan; and [b]Department of Behavioral Science, School of Arts and Letters, Tohoku University, Sendai 980-8576, Japan

Edited by Brian Skyrms, University of California, Irvine, CA, and approved April 29, 2014 (received for review February 14, 2014)

Humans are capable of cooperating with one another even when it is costly and a deviation provides an immediate gain. An important reason is that cooperation is reciprocated or rewarded and deviations are penalized in later stages. For cooperation to be sustainable, not only must rewards and penalties be strong enough but individuals should also have the right incentives to provide rewards and punishments. Codes of conduct with such properties have been studied extensively in game theory (as repeated game equilibria), and the literature on the evolution of cooperation shows how equilibrium behavior might emerge and proliferate in society. We found that community unions, a subclass of labor unions that admits individual affiliations, are ideal to corroborate these theories with reality, because (*i*) their activities are simple and (*ii*) they have a structure that closely resembles a theoretical model, the overlapping generations repeated game. A detailed case study of a community union revealed a possible equilibrium that can function under the very limited observability in the union. The equilibrium code of conduct appears to be a natural focal point based on simple heuristic reasoning. The union we studied was created out of necessity for cooperation, without knowing or anticipating how cooperation might be sustained. The union has successfully resolved about 3,000 labor disputes and created a number of offspring.

reciprocal altruism | imperfect monitoring | proximate cause of behavior

Humans have a unique ability to cooperate, even when cooperation is costly and defection provides a higher payoff. An important reason why humans cooperate in such a situation is that cooperative behavior is rewarded and/or deviations are penalized in later stages. Humans have a cognitive capacity to anticipate rewards and penalties in the future, and they can cooperate by striking a balance between short-run gains and long-run consequences. This basic mechanism to sustain cooperation is commonly referred to as reciprocal altruism in the literature of evolutionary biology and evolutionary psychology, and as repeated game equilibria in game theory and economics.

This observation leads us to the following two fundamental questions to understand the nature of human cooperation:

i) How are rewards and penalties provided to sustain cooperation?
ii) How do cooperation and the system of rewards and penalties emerge?

A substantial body of existing literature has adopted the "top-down" (i.e., theoretical) approach to answer those questions. For example, the theory of repeated games (1, 2) provided theoretical answers to the first question. The theory presents a variety of "equilibria," which specify not only how players should cooperate but also how they should reward and penalize other players. An equilibrium represents stable codes of conduct in the sense that no one has an incentive to deviate from the prescribed behavior. For the second question, the extensive literature on evolutionary game dynamics (3), learning models (4), and simulations/tournaments (5, 6) provided a variety of possible theoretical reasons as to how cooperation may come into existence.

Those top-down works revealed a variety of theoretical possibilities, but we have few clues about which one might be relevant in real-life instances. In contrast, there has also been the "bottom-up" (i.e., empirical) approach to answer questions *i* and *ii*. For example, Elinor Ostrom (7) has conducted extensive field studies to examine how people cooperated in managing "common-pool resources," such as fisheries, groundwater basins, and irrigation systems. The existing empirical works have provided a variety of stylized facts about how people cooperate and how cooperation emerged. Those empirical works, however, are usually not quite informative enough to test the validity of various theoretical possibilities that the top-down literature has identified. Empirical works quite often do not ask key questions that theorists would like to ask. To overcome this problem, the authors, an economist/game theorist (M.K.) and a sociologist (S.O.), have conducted a detailed case study of a labor union. The present paper is a summary of our work in progress.

Overlapping Generations Repeated Game and Community Union

To corroborate theories with the reality, it is desirable that a case study has the following two properties. First, the case should fit the basic structure of a theoretical model well. Second, the case should be simple. In our case study, we chose the overlapping generations (OLG) repeated game as a theoretical model and the community union as a possible real-life instance, because they satisfy these two criteria well.

An OLG repeated game is a dynamic game with OLG of players. Fig. 1 illustrates an example of the OLG repeated game. Each segment of the diagram represents a single player; therefore, at each moment of time, three players are interacting with each other.

Firms and other organizations typically have a structure that is similar to the one described in Fig. 1: The organization itself is an ongoing entity, but its members are short-lived and old members are replaced with young ones. Fig. 1 demonstrates a typical way to sustain cooperation in such a situation. The focal player (yellow) is willing to help a retiring player (gray), because doing so (not doing so) is going to be rewarded (penalized) by a future generation (white). Game theory papers (8–10) showed that cooperation can be sustained as equilibria in the OLG repeated game by means of the chain of rewards and penalties spanning over generations.

One may expect that any company, for example, General Motors, could possibly be the subject of our case study, because it has the OLG structure and workers do cooperate in various ways. The case of General Motors, however, is too complex to examine the kind of theoretical questions we would like to ask about the nature and origin of cooperation. We need a much simpler instance of the OLG repeated game. One of the authors (S.O.) (11) noticed that unique labor unions in Japan, called community unions, are ideal to corroborate the OLG repeated game theory with reality.

This paper results from the Arthur M. Sackler Colloquium of the National Academy of Sciences, "In the Light of Evolution VIII: Darwinian Thinking in the Social Sciences," held January 10–11, 2014, at the Arnold and Mabel Beckman Center of the National Academies of Sciences and Engineering in Irvine, CA. The complete program and audio files of most presentations are available on the NAS website at www.nasonline.org/ILE-Darwinian-Thinking.

Author contributions: M.K. and S.O. designed research; M.K. and S.O. performed research; M.K. and S.O. analyzed data; and M.K. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

[1]To whom correspondence should be addressed. Email: kandori@e.u-tokyo.ac.jp.

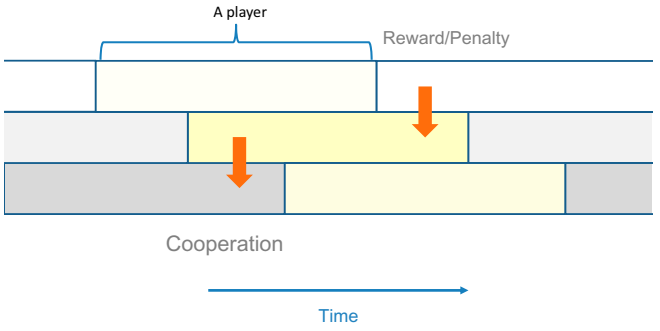

Fig. 1. OLG repeated game.

A community union is an individual-affiliate labor union. When a worker and a firm have a labor dispute about, for example, a layoff, the renewal of an employment contract, or harassment, the worker can obtain membership in a community union even if no other workers from the same firm join the union. The union then exercises its legally protected rights to bargain with and to organize collective actions against the firm, on behalf of that single worker. Community unions are institutions unique to Japan, and they do not exist in the United States and European countries, where individual affiliation to a labor union is not allowed under labor union laws. A substantial majority of labor unions in Japan, however, are firm-affiliate unions, and community unions are a minority. Because there are no official statistics, it is rather hard to pin down the exact number of community unions. The most comprehensive survey of community unions to date (12), however, listed 298 unions as of 2009, and the average membership (based on the 161 unions that responded to the survey questionnaire) was 292. Hence, the total number of community union members can be estimated to be 87,000.

Fig. 2 displays the typical life of a member of a community union. When a worker has a labor dispute with his/her firm, he/she first obtains free telephone counseling services from a community union. If the counseling does not solve the problem, the worker joins the union and the union starts negotiating with the firm on behalf of the worker. Under the labor union law, the firm has a legal obligation to negotiate with the union. If negotiation breaks down, the union goes on to the second stage, where it mobilizes other union members for protest activities, labor relations commissions (LRCs), and litigation. An LRC is a simpler and faster alternative to litigation organized by a local government. When the labor dispute is resolved, the member typically quits the union. As one can see, the community union is a very simple institution whose basic structure closely resembles the OLG repeated game; therefore, it is an ideal subject for our case study.

Case Study: The Tokyo Managers' Union

We focus on one of the larger community unions, the Tokyo Managers' Union. It was founded in 1993 and has resolved about 3,000 cases in the past 16 years (as of 2013; those data were taken from the union's English home page). It has two full-time staff, who are in charge of 231 members (as of 2012). The union was originally created for "managers," such as section chiefs and heads of divisions, but it is currently open to any workers. We collected data directly from the union to verify that it has a structure similar to the OLG repeated game. Fig. 3 shows the inflows and outflows of members in the union's 2012 fiscal year.

The total annual inflow and outflow are approximately balanced, and they are more or less equally distributed over the year. (The union reviews its accounting record at the end of each fiscal year and dismisses members who have stopped paying their membership dues. The larger number of exits in September reflects this practice.) Fig. 4 shows the breakdown of the 108 members who exited in 2012 in terms of how long they had stayed in the union. We were able to collect the data in Figs. 3 and 4 from the accounting record of membership dues.

A substantial majority, 80% of the members (86 of 108), stayed in the union for less than 2 y, and the rest (20%) stayed longer, up to 17 y. Figs. 3 and 4 demonstrate that the membership structure of the Tokyo Managers' Union closely resembles an OLG repeated game.

At this juncture, we must point out that cooperation sustained in the community union is quite striking. The union members come from different firms and do not know each other very well. The turnover rate is quite high, and the members stay in the union for less than 2 y. Reciprocal altruism and repeated game theory have been mainly associated with a long-term relationship between individuals who know each other well in a closed community. The current study is one of the first to show that the same mechanism can operate in a fairly loosely knit organization.

Fig. 5 summarizes the main union activities. Each entry shows the number of cases. For example, in 2011, union–firm negotiation was held for 95 cases (labor disputes). For each case, negotiation typically lasts for multiple stages, so that the total number of negotiation meetings is greater than 95. Unfortunately, the union does not keep records of the number of negotiation meetings. The same caveat applies to other activities.

In about one-half of the cases of negotiation, only the relevant member and staff are involved without any help from other members. Union members mainly help others by attending protest activities, litigation, and LRC meetings.

Fig. 5 shows how many members typically attended each activity. Because the union does not record the number of members attending those activities, the numbers in Fig. 5 are based on our interviews with the two full-time staff members and union members. These numbers show that members are cooperating or helping each other, as the OLG repeated game predicts.

According to the full-time staff, the benefit of cooperation is large: When a large number of members show up, it provides strong pressure on the firm. In addition, for litigation and LRC meetings, the juries and the mediator pay much attention to public opinion; as a result, demonstrating to these entities that the case under consideration has much support has a strong influence on their judgments. The union also regularly invites the media to participate, and cases with a large number of helpers who better attract their attention are likely to be covered.

Members are notified via email of the schedule of planned activities. Aside from protest activities jointly organized with other community unions (whose schedules are sent by the staff), each member is responsible for sending out email notifications to inform others of the venues and dates of his/her own upcoming activities.

How and Why Members Helped Each Other: Other-Regarding Motives?

As one can see from Fig. 5, union members are helping each other by attending other members' protest activities, litigation,

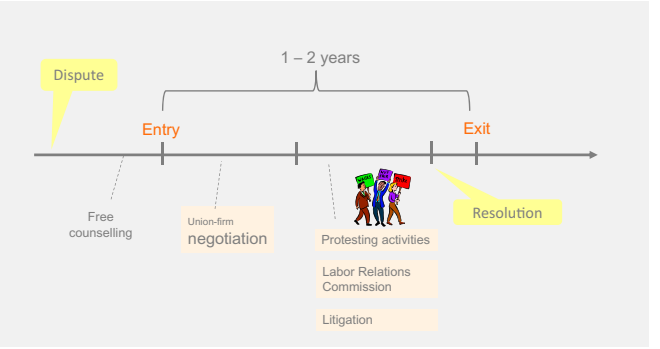

Fig. 2. Typical life of a union member.

Fiscal year 2012	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Total
Entry	9	8	7	13	11	11	15	11	5	6	3	6	105
Exit	3	6	7	4	4	9	2	9	4	6	6	48	108

Fig. 3. Inflows and outflows of members.

and LRC meetings. We tried to uncover the mechanism sustaining cooperation among the members.

First of all, we must consider the possibility that members are helping each other because they are purely other-regarding. We found that other-regarding motives are unlikely to be the main reason for cooperation because of the following considerations. First, cooperation is quite costly. To attend a union activity, a member must take a day off at his/her workplace. This is financially costly, and it also worsens the relationship of the worker with the firm, which has already been strained by the labor dispute between them. Second, Fig. 4 indicates that a substantial majority (80%) of members quit after their cases were resolved. If the members had been purely other-regarding, they would have stayed in the union much longer. When a member is trying to resolve his/her own case, getting help from others is quite important. Hence, a more plausible explanation is that members cooperate in those periods because helping others is reciprocated by some mechanism. (An alternative explanation is that other-regarding behavior is activated only when members are dealing with their own problems. This implies that they cooperate unconditionally in those periods. In what follows, we provide some evidence to the contrary.)

One of the well-documented stylized facts in experimental studies in economics and psychology is that the population is mixed in terms of other-regarding motives. The existing laboratory experiments found that some subjects were purely other-regarding, whereas others were motivated by their private interests. From Fig. 4, we can infer that this is also true in our case but that the fraction of purely other-regarding members appears to be small. One way to measure the other-regarding motive of a member is to examine if she/he stays in the union after her/his own case is resolved. Recall that a substantial majority (80%) of members quit within 2 y, and this indicates that labor disputes are typically resolved within 2 y. Hence, members who stayed in the union for more than 2 y are likely to be the ones who continued to stay in the union after the resolution of their own cases. This is a conservative measure, because it might take more than 2 y to resolve a dispute. The fraction of purely other-regarding members thus measured is ∼20%, according to Fig. 4. A substantial majority (80%), in contrast, quit (and thus stop cooperating) right after the resolution of their own cases. These members might be other-regarding to some extent, but their other-regarding motive is not strong enough to make them stay in the union (and help others) after the resolution of their own cases.

Given those observations, we proceed to examine the possibility that the members cooperate because cooperation is reciprocated by some mechanism. Let us first describe theoretical possibilities identified by the existing literature in game theory.

Equilibria in the OLG Repeated Game and the Problem of Observability

The literature in game theory and economics (8–10) has identified a variety of equilibria to sustain cooperation in the OLG repeated game. An "equilibrium strategy" specifies how a player should adjust his/her behavior based on past history in the game. Moreover, it has the property that the player does not have any incentive to deviate from his/her strategy after any history. This means that if cooperation is sustained in equilibrium, not only are players willing to cooperate but they also have the right incentive to reward honest players and to punish deviators. Hence, a repeated game equilibrium is a candidate for a sustainable code of conduct for cooperation. If cooperation is based on a nonequilibrium strategy, a profitable deviation is found sooner or later and cooperation is likely to break down.

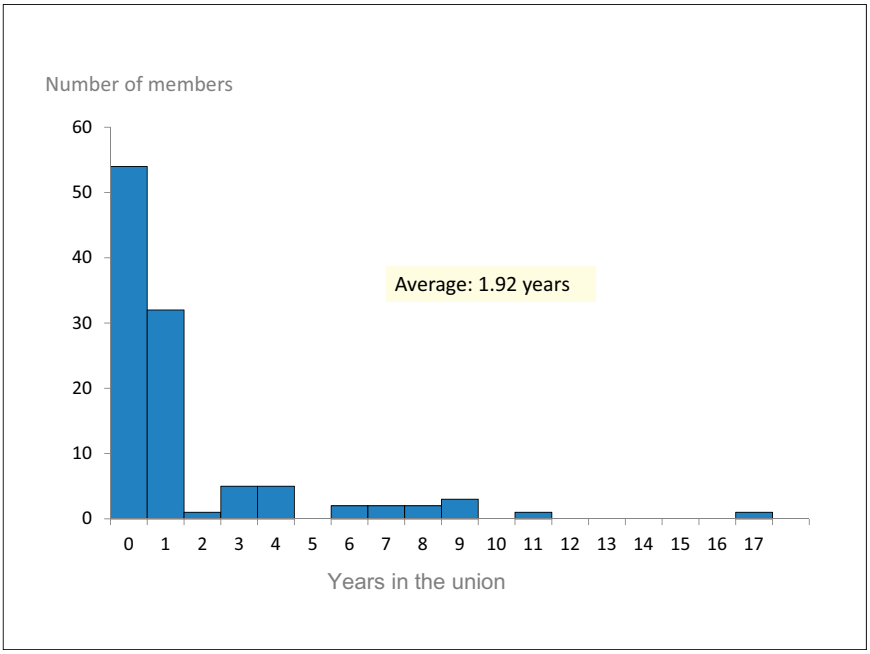

Fig. 4. Years in the union.

	How many members are typically mobilized ?			
		5-20	10	5-20
	Negotiation	Protesting	Litigation	Labor Relations Commission
2008	186	10	16	13
2009	128	23	9	20
2010	169	21	10	15
2011	95	19	10	24

Fig. 5. Union activities and cooperation. The numbers of mobilized members are based on interviews with the full-time staff. The remaining numbers were collected from minutes of annual meetings of the Tokyo Managers' Union.

The existing game theory literature on the OLG repeated game typically assumes perfect monitoring: At each moment in time, players can observe the entire history of past actions. We found that this assumption is by no means satisfied in our case, and one of the challenges we face is to identify or construct an equilibrium that can work under the very limited observability of members' behavior in the union. Before describing the equilibrium we found, let us review the common properties of the equilibrium in the existing literature that assumes perfect monitoring.

A simple, and probably the best-known, equilibrium is based on the "trigger strategy." It specifies that players should cooperate (except for the last few stages before retirement, when they do not face significant rewards and punishment in the future). If anyone deviates, then all players stop cooperating forever. Players are willing to cooperate because deviation destroys the benefit of cooperation in the future. Moreover, players are willing to punish (i.e., to stop cooperating with) a deviator because they do not have any incentive to cooperate after a defection (no one else ever cooperates again). This is, however, a fragile equilibrium, where cooperation among all members breaks down after a single episode of defection. We found that this is highly implausible in our case, and we found no support for the trigger strategy equilibrium in our interviews with the union staff and the members.

The problematic feature of the trigger strategy equilibrium is that all players are punished after the defection of a single player. A more realistic equilibrium would be based on personalized punishment: If a player helps others, he/she can obtain help in the future; otherwise, no one helps that player. A previous paper written by one of the authors (S.O.) constructed such an equilibrium. In such an equilibrium, if a player is not helping someone, it might possibly mean that (*i*) he/she is defecting or (*ii*) he/she is punishing a defector.

In the former case, the player should be punished; however, in the latter case, the player should not be punished (otherwise, he/she has no incentive to punish a defector). To distinguish between cases *i* and *ii* above, the union members must have somewhat detailed knowledge about what has happened in the union. Do they actually have the necessary information? We now focus on this question.

Observability and Reputations in the Union

We found that observability of union members' activities is quite limited. The union does not keep records of who helped whom. (The union started to post sign-up sheets of union activities in the union office in December 2013. They also used to utilize sign-up sheets before the summer of 2010.) Not a single member we interviewed had that information. The assumption of perfect monitoring, which is widely adopted in the OLG repeated game literature, definitely does not hold in the Tokyo Managers' Union.

There is, however, a possibility that the equilibrium under perfect monitoring can be implemented by the reputation mechanism. For example, a member may carry a reputation as a "good citizen" or as a "free rider." A good citizen does not lose her reputation when she refuses to help a free rider, but her reputation is lost when she does not help another good citizen. In this way, reputation can summarize the relevant information about past history and players can effectively implement an equilibrium strategy (i.e., they can distinguish between cases *i* and *ii* in the previous section) even if they only observe their reputations, rather than the whole history of actions in the union. The theoretical possibility of a reputation mechanism, whereby players carry reputations that are systematically revised, has been explored by game theory literature (13) and evolutionary biology literature (14) (in the latter, the reputation mechanism is called "image scoring").

Our interviews with the staff and members, however, did not confirm the existence of the reputation mechanism in the union. First, members do not have any information about what has happened in the union when they join. Second, there are limited opportunities for the members to exchange information about the members' reputations. There is no opportunity for all members (~300) to congregate. The number of key union activities and their attendance can be seen in Fig. 5, and in addition to those activities, the union organizes other events where members can socialize and may potentially exchange information about reputations (Fig. 6).

Recall that the members are typically from different firms and that they do not come to the union office unless they have something to do in the office. The activities in Fig. 6 appear to be insufficient for the members to share each other's reputation. The most notable opportunity for the members to know each other is the Managers' Action Committee (MAC), which refers to monthly or bimonthly meetings where union members exchange information about their problems and upcoming activities (negotiation, protest activities, litigation, and LRC). Twenty-four members attended the MAC meeting on December 11, 2013. They introduced themselves and explained their problems. The chairman of the meeting was a full-time staff member, and he provided useful information and advice. Seventeen members were at the stage of negotiation, three were appealing to the LRC, one was filing a lawsuit, one was organizing protest activities, and two had already resolved their disputes (two members were engaged in multiple activities). The members asked for help, but no exchange of information about the reputations of members was observed (including the conversations before and after the meeting). This meeting also gave an impression that the staff had fairly detailed knowledge about each member's profile but that the members did not know each other very well.

We specifically asked the staff and members if they had heard about the reputations of other members. The answers were unanimous. They reported that they had not heard of (good or bad) reputations in the context of how much help a member provided to others. Some of them mentioned that they had heard about members' bad reputations but that they were associated

	2010	2011	Number of attendance per meeting
Orientation meeting	2	2	10–15
Study group	18	12	15–20
Resolution party	2	4	40
Recreational activities	2	5	10–40
Managers' Action Committee	12	6	20–30
Annual Meeting	1	1	100

Fig. 6. Opportunities for socialization. The numbers of attendance are partly based on interviews with the full-time staff. The numbers were collected from minutes of annual meetings of the Tokyo Managers' Union.

with problematic personalities, with soliciting contributions to a dubious investment project, or with a pyramid scheme.

Possible Equilibrium in the Union

Hence, the well-known equilibria identified in the existing literature, which are based on perfect monitoring or reputation, do not seem to explain cooperation in the Tokyo Managers' Union. To find out how cooperation is sustained, we looked for union members who had recently resolved their cases and could tell us about their experiences. This was a difficult task because the members typically quit after their cases are settled, but we found one member who agreed to have an interview with us.

The interview was conducted on October 29, 2013. "Mr. A" (pseudonym) has been a union member since 2011. He had three LRC meetings, after which his case was resolved in July 2012. At the end of this section, we will present an equilibrium that is consistent with this interview. We did not ask any leading questions to support this equilibrium, simply because we were not aware of this equilibrium at the time of this interview. The equilibrium was constructed after the interview, based on Mr. A's experience and view.

Basically, we let Mr. A tell us his own experience and opinions. We asked questions in the following order:

i) The nature of his labor dispute and how he joined the union.
ii) How his union–firm negotiation was organized.
iii) How his LRC meetings were organized. We asked how many other members came to help him and how he met them.
iv) How much help he provided to others. We asked how many times he attended other members' protest activities, LRC meetings, and litigations. We also asked how he met those whom he helped.
v) Whether he has heard about members' reputations regarding how much help they had provided to others.
vi) Miscellaneous questions (what he had done after the resolution of his own case, how long it typically takes to resolve a member's case, etc.).

Item *v* is the only question that is directly related to a theoretical hypothesis. Its purpose was to examine if the reputation mechanism, often postulated in the existing literature, is present in the union. Originally, we suspected that this would be the mechanism to sustain cooperation in the union. As we have already described in the last section, however, Mr. A's answer was negative.

Let us now describe the contents of the interview. How much help did he get? A full-time union staff member, a person who had been a union member for a long time, and a couple of other members attended his LRC meetings. Mr. A did not recall their names, but he came to know them when attending activities to help other union members.

Mr. A helped other members extensively. He attended approximately 20 protest activities, 10 LRC meetings, and 5 trials. The LRC meetings he attended concerned disputes of different members. He said that he came to know those members "when distributing leaflets at protest activities (for other members)," and they asked Mr. A to come to their LRC meetings. Mr. A decided to carry a datebook to take note of the schedule of those meetings.

When asked why he attended those activities, Mr. A replied, "That was because I wanted them to come to my LRC meetings and trials. If you want them to come, you must also go." He told us that he tried hard to be visible. He lived in a suburban area far away from the central part of Tokyo, but he chose to ride a bicycle to attend union activities. "When I said that I came all the way from [his town] by bike, wearing a helmet and biking outfit, I was very visible. That was very effective." Mr. A had not attended any of the MAC meetings.

We examined if there is any way to formulate Mr. A's story as an equilibrium of an OLG repeated game theory with limited observability. The essential feature of Mr. A's story is that when he came to help someone, he met other members and they asked Mr. A to help them. The simplest model we found to formulate Mr. A's experience and view is shown in Fig. 7.

We found that the following strategy, which is consistent with Mr. A's view, constitutes an equilibrium:

i) A young player helps a randomly selected player. This corresponds to the fact that Mr. A had helped a large number of players before he had his own LRC meetings.
ii) A middle-aged player helps a player he met (when helping someone when she/he was young). This formulates the essence of Mr. A's story: Members who come to help the same person will help each other.
iii) Old players do not help others. This is consistent with the fact that the union members typically quit (and no longer help others) after their cases are resolved.

Let us check the key incentive constraints.

a) Incentive to cooperate: A player has an incentive to help another player, because that gives him an opportunity to meet other members who are going to help him.
b) Incentive to provide rewards and punishment: Why does a player help (reward) the "right" person? Assume that players i and k came to help player j. In the equilibrium, player i will help player k in the next period, thereby rewarding player k for helping player j. Why does player i help player k and not another player? Note that helping others is profitable (as stated in *a*) but that players do not have enough time to help all players. Therefore, player i must choose which player to help. This is the key to understanding our equilibrium. In the equilibrium, player i is basically indifferent about whom to help; the equilibrium treats players symmetrically, and therefore in terms of meeting others (that is the sole reason in the model to help someone as pointed out in *a*), helping any player is equally worthy. As a result, helping someone whom a player met (i.e., providing a reward to player k) is one of the optimal actions of player i.

The explanation above provides the essence of the equilibrium, and we now provide more details. The properties *i–iii* do

- The model has an infinite time horizon t=1,2,3,...
- In each period, two players enter the union.
- Each player stays in the union for three periods. A player is called *young*, *middle-aged* and *old* respectively, in her/his first, second and third periods in the union. Hence, in each period there are six players (two young, two middle-aged, and two old players.)
- A player can help at most one player in each period.
- Players need help when they are middle-aged and old. The current members are informed of the names of those players who are seeking help (via e-mail).

Observability: Young players do not know what has happened in the union before they enter, nor do they know the ages of the existing members. When a player helps someone, she/he observes who else also came to help the same person. This is the only information a player has about other players' behavior.

Payoffs: The benefit and cost of one help are denoted by b and c. Player i's payoff is equal to

$$(\text{total number of helps i obtains}) \times b - (\text{total number of helps i provides}) \times c$$

Fig. 7. Model of a community union.

not completely describe a strategy. For example, if a middle-aged player met more than one player, which one should she/he help? What if she/he met no one? To obtain the exact equilibrium conditions, those fine details should be specified in a proper way. We do not believe that the union members exactly follow those details, nor do we expect them to calculate the cost and benefit of cooperation very precisely. Our point is that the basic structure of our equilibrium, *i*–*iii*, matches Mr. A's experience and view and that the essence of the incentive constraints, *a* and *b*, can possibly be understood by the union members. The rough sketch of the code of conduct presented above can be made into a precise game theoretical equilibrium by, for example, the following elaborations on *ii*:

iia) If a middle-aged player met more than one player, she/he helps the one in the same generation (i.e., a player who entered the union in the same period as the player did) if there is such a player. Otherwise, she/he randomly chooses whom to help.

iib) If a middle-aged player did not meet anyone, she/he helps the other player in the same generation.

With this specification, one can calculate the cost and benefit to follow the equilibrium behavior in each period. Checking those incentive constraints reveals that the strategy above is an equilibrium if $b > 8c$, where b is the benefit of a single instance of help and c is the cost of a single instance of help.

How does this equilibrium overcome the difficulty associated with the limited observability in the union? First, note that there is always a possibility that more than one member may come to help someone and that they mutually monitor each other's cooperative behavior. The reward for cooperation is provided by a member who directly observed it. Second, in this equilibrium, no one needs to monitor if a member is rewarding the right person (this is the key feature of this equilibrium). The equilibrium makes helping anyone more or less equally profitable; therefore, helping the right person is one of the best actions.

We must stress that it would have been virtually impossible for us to come up with such an equilibrium had we not had an interview with Mr. A. This equilibrium formulates how Mr. A envisioned the code of conduct in the union, but we have yet to confirm that the same view is widely shared by other members. Hence, it is premature to claim that all union members are actually following this equilibrium. We plan to continue our study to collect further information to uncover the mechanics of cooperation inside the union.*

Episode of Deviation

We then looked for some supporting evidence about the mechanics of cooperation described in the previous section. The point is that the members are not purely other-regarding players who help others unconditionally. Instead, the members are conditional cooperators, and failing to help others would result in no help from others. Is there any episode where a member got no help from others? A full-time staff member and a founder of the union, Kiyotsugu Shitara, described an illuminating episode in an interview conducted on October 6, 2013. (In the interview, we explicitly asked if he had known any episode where no one came to help a member. We did not tell him about our equilibrium hypothesis.) During a protest activity in 1995, Shitara and another member of the union found themselves alone. No one came to help, and the member, for whom the protest activity was organized, almost burst into tears. The reason why no one showed up was that Shitara had forgotten to suggest to him that he should ask other members in a face-to-face conversation to help him. (This is an important piece of advice that he normally gives to all members.) The member in question actually had interacted with others, but he had failed to ask them for a help. Kiyotsugu Shitara smiled and remarked, "That was terrific. All members should have such an experience." He believes that such an experience would work as good "training" for a member, to make him/her realize how one should try hard to get help from others.

Emergence of Cooperation in the Union

We now examine various theories of the evolution of cooperation, given the reality of the community unions. First, we examine how members come to establish an equilibrium to sustain cooperation within a union. In the next section, we examine the evolution of community unions.

A number of theoretical possibilities have been suggested in the existing literature about how equilibrium (to sustain cooperation) might emerge:

Social contract: Members explicitly discuss and agree on how to cooperate.

Dynamic selection: Equilibrium to sustain cooperation emerges by the survival or learning of successful strategies.

Social norms: There is a publicly known code of conduct about what the individuals should do.

Focal point: The equilibrium to sustain cooperation "stands out" as a natural thing to do.

The first (social contract) might be a good explanation for cartel or bidding rings (where individuals explicitly discuss how to collude), but it obviously does not apply to community unions, where members exit and enter frequently.

Concerning the social norms explanation, we were rather surprised to find out that in the Tokyo Managers' Union, there is virtually no explicitly stated code of conduct or rules about how members should help each other. The organization's home page explains the philosophy of the union as follows:

> . . .We stress that the union does not tell you what to do or bail you out from your trouble. It is up to you to resolve your own case. The union listens to your problems and supports you, but you must take an initiative to resolve your case. . . It is other union members who [help you], and they are dealing with their own problems. The union is based on the voluntary activities of its members. You are free to attend any meetings, to say whatever you like, and to do whatever you wish, but you are responsible for the consequences. The Tokyo Managers' Union is a free and voluntary organization for someone like you.

In his interview, Kiyotsugu Shitara, a full-time staff member, elaborated on this point. He stated that acting together with other members and asking for help in person are essential, and that he always tells this to the members. Not all members, however, listen to him, but he embraces it. He said, "If you are too picky about the degree of contribution [of a member] to the union or the frequency of participation, things do not work out. . . Everything should be voluntary."

Mr. A told us in his interview that he helped others because he thought that would be the only way to get help. He reported that no one told him to do so, but he thought about Kiyotsugu Shitara's statement that "the union does not tell you what to do or bail you out from your trouble" and reached the conclusion that helping others is the right thing to do.

Hence, the mechanism under which members come to cooperate appears to be somewhere in between the social norm and the focal point. The evolution/learning mechanism seems to play little role in this case. Mr. A, without any trial and error, immediately came up with the right thing to do. The equilibrium we described in the previous section naturally came to Mr. A's mind as a focal point. The union also suggests how members

*An alternative explanation of cooperation in the union is that members attend protest activities, LRC meetings, and litigations for other members because they provide an opportunity to learn how to resolve labor disputes. This explanation does not seem to apply to Mr. A, because the number of meetings that he attended ($n = 35$) appears to be more than enough for learning. Also, he explicitly stated that he helped others because that was the only way to get help. We have yet to confirm, however, that the same is true for other union members.

should help each other, but not so explicitly and forcefully. The sense of voluntary cooperation appears to be fairly important.

Evolution of Community Unions

Community unions emerged around 1980. Individual affiliation to a labor union had been technically possible previously, but the existing unions were reluctant to grant individually based membership. Around 1980, the number of labor disputes about layoffs sharply increased, after the crash of bubbles in Japanese financial markets. A lawyer specializing in labor disputes recalls that he asked some labor unions to accept individual affiliations to deal with the increasing number of labor disputes but that they were very reluctant (page 6 of an essay by Kenji Tokuzumi in the pamphlet for the 15th anniversary of the Tokyo Managers' Union in 2008). The reason he heard was that accommodating individual labor disputes does not contribute to organizing the union; it takes a substantial effort and cost to help individuals, but they would quit after their problems are resolved.

Hence, organizing workers from different firms for union activities had been deemed impossible. How did community unions, which challenged this conventional wisdom, emerge? In the theoretical model about the evolution of cooperation, a "mutant" strategy is introduced out of the blue, but how did it happen in reality? In our interview with Kiyotsugu Shitara, one of the founders of the Tokyo Managers' Union, he explained his own experience in creating the union back in 1993. Basically, he created the union out of necessity. The Japanese lifetime employment system was falling apart, and there were an increased number of individual labor disputes concerning managers' layoffs. As a leader of a labor union confederation in Tokyo, he organized free telephone consulting for troubled workers and found out that the problem was quite widespread. He had a vivid memory about the very moment when the idea of his union emerged. He had a conversation with a newspaper reporter and a television reporter who came to cover the telephone consulting he was organizing. They reached a conclusion that something more systematic than telephone consulting should be done, and Shitara told himself that a labor union for managers might be a solution. The reporters were all excited to hear what he said and told him, "That's it! That is what you should do."

Kiyotsugu Shitara told us that he had no idea about how to organize union members, who come from different firms and tend to quit after their cases are resolved. He stated that helping those troubled workers was most important to him and how to do it was of secondary importance. He stated, however, that from the very beginning, the members helped each other. When asked why cooperation naturally emerged, he told us that back then, the Tokyo Managers' Union had only one full-time staff member and the union could not possibly function without the mutual help of the members.

In summary, cooperation between short-lived union members, which had been deemed impossible, emerged out of necessity. From the very beginning, cooperation was sustained, possibly because the structure of the union had the right characteristics to support a cooperative equilibrium, which was a natural focal point.

Lastly, let us examine how the successful strategy (the system of cooperation in the Tokyo Managers' Union) proliferates in society. In the theoretical literature of evolutionary game theory, a successful strategy has a large number of offspring (by means of reproduction in biological applications and by imitation in social or economic applications). Let us examine how this applies to the case of community unions. Did the Tokyo Managers' Union imitate a (successful) existing strategy? Kiyotsugu Shitara told us that when he created his union, he only had a basic idea about what other community unions had been generally doing, but he said that he did not study how other unions organize their activities or copy their practices. In summary, Shitara did not imitate other unions' practices, but he knew that other unions were functioning well. The Tokyo Managers' Union is doing well and creating offspring organizations. Fig. 8 summarizes the spinoffs from the union. (Those are the unions that Shitara considers to be the union's offspring. Information on founding years, terminating years, and membership was collected by interviews with the relevant parties.) The thick lines in Fig. 8 indicate the cases where the Tokyo Manager's Union provided financial and human resources in the first year of operation.

Concluding Remarks

Let us summarize the main messages of our study. We found that community unions, a subset of labor unions that admits individual affiliation, have a structure that closely resembles the OLG repeated game theory. A case study of the Tokyo Managers' Union revealed that members cooperate, even if they enter and exit frequently and do not know each other well. The existing literature on

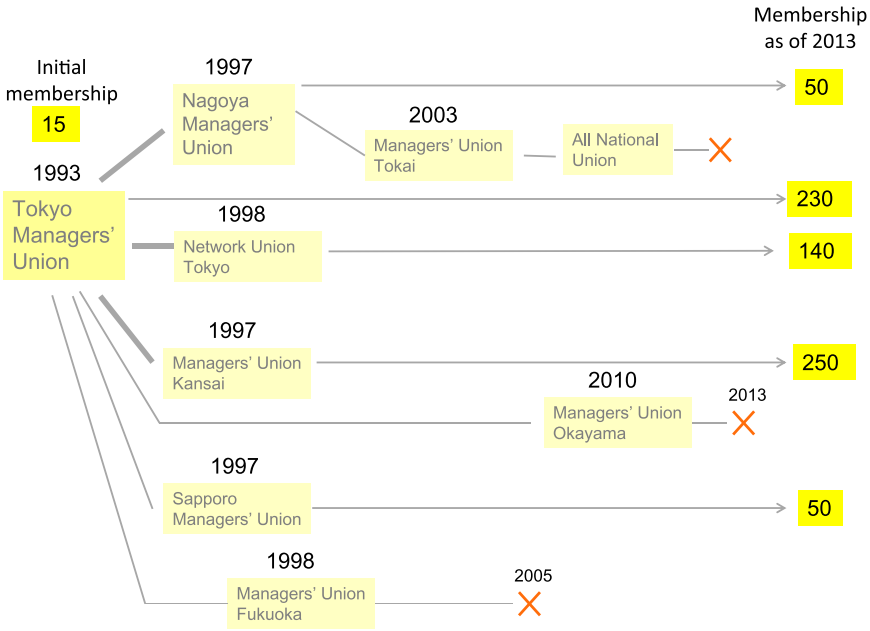

Fig. 8. Offspring of the Tokyo Managers' Union.

OLG repeated game theory has shown that cooperation would be possible in the union, assuming that (*i*) players can perfectly monitor everyone's past actions or (*ii*) players carry reputations that summarize their past experiences. Our case study shows that neither is true in the union. Based on a detailed interview with a member of the union, we discovered an equilibrium that can sustain cooperation under the very limited observability in the union. We also examined the evolution of the union (i.e., how the union was created and how it has produced offspring). A notable insight we obtained is that cooperation can be sustained in a very loosely knit organization where members, who do not know each other well, enter and exit quite frequently.

Our maintained hypothesis is that sustainable cooperation is likely to be a game theoretical equilibrium, which admits no profitable deviations. This means that when individuals have a short-term incentive to defect from cooperation, there must be a system of rewards and/or punishment. We believe that understanding this mechanism, which determines how cooperation is reciprocated, is important in understanding how cooperation is sustained. There are, however, a substantial number of individuals who are antagonistic to conscious and rational calculation of short-term gains and long-term losses, and they dismiss such a practice as being quite selfish and antisocial. Does the presence of those people imply that game theory is not useful? We argue to the contrary. If cooperation is not sustained by an equilibrium code of conduct, a profitable deviation is found (most likely by chance, rather than by rational calculation) sooner or later and cooperation is likely to erode. This is the main reason why we used a game theoretical analysis to examine sustainable cooperation.

Our study suggests the value of case studies to understand human cooperation. Case studies are a common mode of research in management sciences, but they are not a mainstream methodology in the literature on repeated game theory. Note that a repeated game has a large number of equilibria. [The folk theorem for repeated games (1, 2) shows that a repeated game typically has infinitely many equilibria if players are patient.] Some of them are well known, such as the trigger strategy or tit-for-tat, but not all equilibria have been fully identified. One revealing lesson we learned in our case study is that every case has a unique set of details, and it is highly unlikely that a well-known equilibrium works as it is. The challenge is to identify or construct an equilibrium that fits the "local climate" of the case under consideration. This suggests the following research strategy, which proceeds in two steps. The first step is finding a case that satisfies the following two criteria: (*i*) it should fit the basic structure of a theoretical model well, and (*ii*) it should be simple enough. This is the most important task. We regard finding the striking similarity between the OLG repeated game and community union as an essential part of our study. The second step is identifying or constructing an equilibrium that fits the local climate (a unique set of details) of the case.

Repeated game theory has not identified all equilibria, but it does provide various techniques to construct equilibria. Combining the working knowledge of repeated game theory and detailed empirical case study, one could identify an equilibrium played in the case under consideration (if any). This is closest in spirit to what the economic historian Avner Greif called "interactive, context-specific analysis." In his renowned monograph (15), he explained it as follows: "Central to this method is a context-specific analysis that interactively uses deductive theory and context-specific modeling to develop and evaluate conjectures about the relevance of particular institutions."

Using this method, Avner Greif constructed a possible repeated game equilibrium that sustained cooperation in long-distance trade in the 11th century by Maghribi traders, based on the documents found in the *geniza* (depository) in Old Cairo. This method can be applied well beyond the study of economic history, and we believe that the case study in this context is a promising way to examine the nature of human cooperation. We must note, however, that a case study has its upsides and downsides. A drawback is that a case study is hard to generalize from because (by definition) it concerns a single observation. A benefit is that it provides fairly detailed and reliable information on the case under consideration. Our hope is that as the number of case studies increases, we can obtain reliable and detailed information about the nature of human cooperation.

Lastly, let us summarize the implications of our study in a broader context of the evolution of cooperation. Biologist Nikolaas Tinbergen advocated an invaluable distinction between the proximate vs. ultimate cause of behavior. The former refers to the physiological or cognitive mechanism to deliver the behavior, and the latter concerns the evolutionary origin of the proximate mechanism.

As for the proximate cause of cooperation, our study suggests that humans are not directly programmed to play a particular repeated game equilibrium strategy (e.g., the grim trigger or tit-for-tat). Our case study suggests that every case has a unique set of details and that the mechanism of reciprocal cooperation (a repeated game equilibrium) should be crafted in such a way that it works given those details. It is highly unlikely that a small set of behavior works in all cases. The same view was echoed by Elinor Ostrom, who conducted extensive empirical studies of human cooperation in the management of common-pool resources (16). We found a specific equilibrium that can work under the limited observability in a community union, and a member (Mr. A) naturally followed this by heuristic reasoning that the only way to get help is to help others. This heuristic is fairly general, but it somehow pinned down a particular mode of behavior that works in the union. Hence, our view in the broader context of human evolution is that humans are likely to have some general heuristics that make some repeated game equilibria focal, and the heuristics are flexible enough to produce a variety of equilibria to sustain cooperation, which fit various local conditions. We have yet to understand the nature of those heuristics and exactly how they work. These heuristics are of central importance to the study of human cooperation. The question about the ultimate cause, which asks how these heuristics have evolved, is a challenging agenda for future research.

ACKNOWLEDGMENTS. We thank Kiyotsugu Shitara, Tsuyoshi Suzuki, Masaaki Koike, and the members of the Tokyo Managers' Union for their cooperation. We also thank Ryo Takahashi for editorial assistance. This study was supported by the Japan Society for the Promotion of Science (JSPS) [Grant-in-Aid for Scientific Research 21530165 (to M.K.) and Grant-in-Aid for JSPS Fellows 24-6997 (to S.O.)].

1. Mailath G, Samuelson L (2006) *Repeated Games and Reputations* (Oxford Univ Press, New York).
2. Kandori M (2008) Repeated games. *The New Palgrave Dictionary of Economics*, eds Durlauf S, Blume L (Macmillan, New York), 2nd Ed.
3. Weibull J (1995) *Evolutionary Game Theory* (MIT Press, Cambridge, MA).
4. Kalai E, Lehrer E (1993) Rational learning leads to Nash equilibrium. *Econometrica* 61:1019–1045.
5. Nowak M, May R (1992) Evolutionary games and special chaos. *Nature* 359(6397): 826–829.
6. Axelrod R (1982) *The Evolution of Cooperation* (Basic Books, New York).
7. Ostrom E (1990) *Governing the Commons* (Cambridge Univ Press, New York).
8. Cremer J (1986) Cooperation in ongoing organizations. *Q J Econ* 101(1):33–49.
9. Kandori M (1992) Repeated games played by overlapping generations of players. *Rev Econ Stud* 59(1):81–92.
10. Smith L (1992) Folk theorems in overlapping generations games. *Games Econ Behav* 4(3):426–449.
11. Obayashi S (2013) The mechanism of maintaining a reciprocal system in a group with flexible membership: an empirical and mathematical approach. *Shakaigaku Hyoron* 64(2):240–256.
12. Suzuki R (2010) [*The Report on a Questionnaire Survey on Individual-Affiliate Labor Unions*] (The Ohara Institute for Social Research, Hosei University, Tokyo), Working Paper No. 41. Japanese.
13. Kandori M (1992) Social norms and community enforcement. *Rev Econ Stud* 59(1):61–80.
14. Nowak MA, Sigmund K (1998) Evolution of indirect reciprocity by image scoring. *Nature* 393(6685):573–577.
15. Greif A (2006) *Institutions and the Path to the Modern Economy* (Cambridge Univ Press, New York).
16. Ostrom E (2010) Beyond markets and states: polycentric governance of complex economic systems. *Am Econ Rev* 100(3):641–672.