Chapter 12

Measuring Empowerment at the Community Level: An Economist's Perspective

Asim Ijaz Khwaja

Experiences over the past few decades suggest a shortcoming of top-down approaches to development. Since the 1980s, the new watchwords have been "participatory" or "community-led" development (Mansuri and Rao 2004; Uphoff 1996) and, more recently, "empowerment." The World Bank's *Empowerment and Poverty Reduction: A Sourcebook* defines empowerment as "the expansion of assets and capabilities of poor people to participate in, negotiate with, influence, control, and hold accountable institutions that affect their lives" (Narayan 2002, 14). Before empowerment can be integrated into development policy, however, it must be clearly conceptualized, and reliable measures must be developed. This is particularly important given that such measures of empowerment are likely to become project goals for development agencies.

This chapter offers an economist's perspective on how one would begin to construct measures of empowerment and the issues involved in doing so. I do not propose to offer a laundry list of potential measures applicable in all circumstances; such an exercise is almost certainly futile, as good measures are likely to be context-dependent. A more promising approach is to develop a general framework for conceptualizing empowerment, which can then be employed by researchers and practitioners to develop measures appropriate to a particular context and goals. While the chapter puts forward a theoretical framework, it also draws heavily on my own empirical work (Khwaja 2001, 2004) in order to demonstrate how this framework can be applied in the field and to provide empirical support for it.

Even before developing a general framework, it is imperative to make some basic distinctions among alternative understandings of empowerment. In particular, the researcher/practitioner needs to clarify whether empowerment is construed as an *end in itself* or as a *means to an end*, or both. The chapter accordingly begins with a section on the importance of making this distinction.

The second section develops a particular theory of empowerment and explains how this theory allows us to construct context-specific measures of empowerment. In order to illustrate both the theory and such measures, this chapter considers a specific case based on an empirical study conducted by the author, namely, the collective maintenance of community infrastructure projects in Baltistan, northern Pakistan. The study uses data from a primary survey I conducted of 99 rural communities in the region. In addition to community- and household-based surveys, the study also included technical surveys of all externally initiated infrastructure projects in these communities. The primary purpose was to examine the determinants of a community's collective success in maintaining these public infrastructure projects. Since maintenance of these projects was solely the responsibility of the community, their upkeep also provided a measure of the community's collective potential (for detailed empirical results of the study see Khwaja 2001).

The third section of the chapter describes an important and difficult empirical issue that arises if empowerment is to be viewed as a means to an end: establishing that there is a causal relationship between a particular measure of empowerment and the outcome or end of interest. This section also draws on my previous work to illustrate how one can investigate such causal inferences. More generally, this section cautions that while theory can suggest empowerment measures, causal channels should be subjected to rigorous empirical analysis. This is necessary in order to show that the measures affect outcomes of interest and are consequently useful for policy.

Empowerment: An End or a Means to an End?

The literature on empowerment shows two understandings of the concept. Empowerment is sometimes understood as a means to a specific end, such as increased welfare of the empowered agent. It is also often conceived as an end valuable for its own sake. It is hardly surprising that a concept as broad as empowerment can be understood in more than one way. However, if one is to develop and then use measures of empowerment in policy initiatives, it is essential to be explicit about which understanding is being used. The theoretical framework and measures constructed, as well as the process of establishing whether the measures have a causal impact, will be quite different in the two cases.

For an illustration of these issues, consider one aspect of empowerment as defined in the World Bank's empowerment sourcebook (Narayan 2002):

expanding poor people's capabilities. This is undoubtedly an important goal. Before attempting to develop measures, however, it must be clarified whether the researcher is asserting that such capabilities are important *only* because they lead to an increase in the welfare or well-being of the poor, as measured by standard socioeconomic indicators, or whether an expansion in these capabilities has value even if they do not influence any other aspect of welfare. In either case, the assumption is that empowerment is valuable because it affects an agent's overall welfare. The distinction here is whether this effect is true by definition, that is, empowerment is *defined* as a component of an agent's welfare or utility (empowerment as an end), or whether it is true by causation, that is, empowerment *influences a component* of welfare such as the agent's income or health status (empowerment as a means to an end).

Expressed mathematically, one is trying to posit whether the relationship between a particular aspect of empowerment is given by equation (1) below or by the systems of equations (2), where U_i is an agent i's measure of welfare, E_i is a measure of how empowered she is, X_i is a list of other factors that directly affect her welfare, and $f(.)$ and $g(.)$ are functions. The triple equal sign indicates that the relationship is posited to be definitionally true and must therefore be defended as such. These equations are illustrated in figure 12.1.

(1) Empowerment as an end: $U_i \equiv f(E_i, X_i)$

(2) Empowerment as a means: $U_i \equiv f(X_i)$ and $X_i^j = g(E_i)$ for one or more factors j

Now consider the process of laying out a theoretical framework and constructing measures using the second understanding of empowerment—empowerment as a means to achieve a specific end. If, for example, we view "expanding poor people's capabilities" as a means to achieve greater income of the poor, then we first need to articulate how such expanded capabilities can lead to an increase in income, that is, the nature of function $g(.)$ above. The definition of empowerment used in the sourcebook suggests one potential channel: expanded capabilities allow the poor to "*influence* . . . institutions that affect their lives" (Narayan 2002, 14, emphasis added). That influence, it may be hypothesized, can in turn affect their income.

To take an example from my work on public infrastructure projects, the Baltistan study mentioned above, consider delivery of a local public project to a community. Suppose that the poor in the community would benefit from an irrigation channel. However, the better-off minority in the community prefer an alternative project, such as electricity generation. In the absence of influence by the poor, the latter project will be chosen. If the capabilities of the poor are sufficiently expanded, however, they can exert their influence in favor of the project that provides most benefit to them. Such a theoretical framework would suggest that measures of empowerment should reflect capabilities related to such potential influence. One measure could be whether the poor were given an equal vote in project selection, for example, by the project provider going to all community households and asking them which project

FIGURE 12.1 Empowerment as an End or as a Means to an End

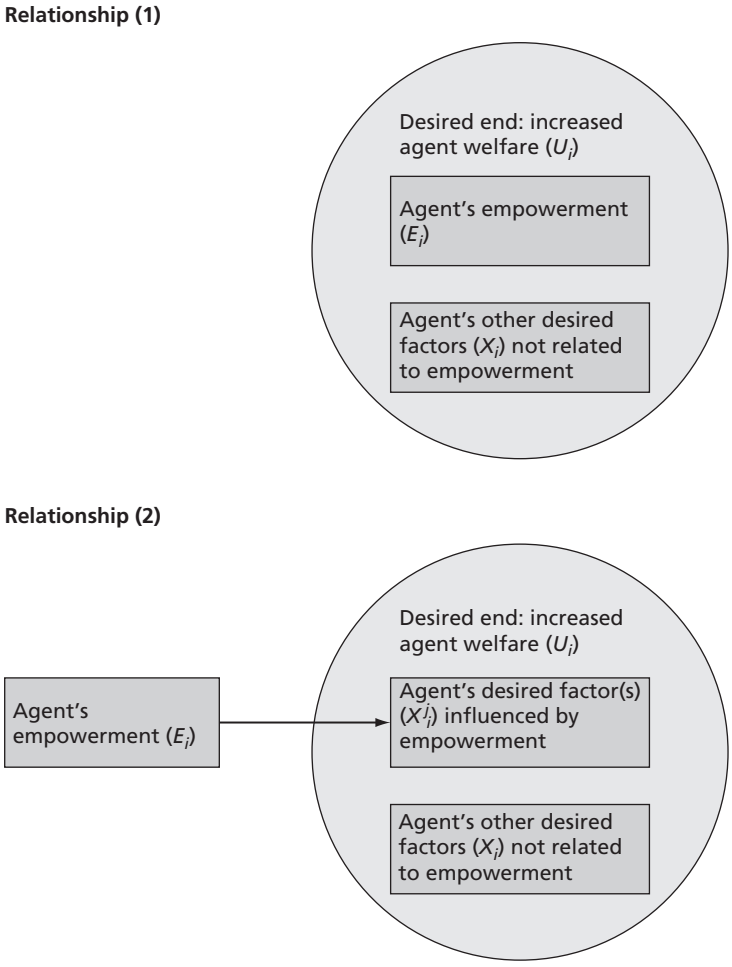

they prefer. In another context we may arrive at a different measure, based on another possible causal relationship. The point is that specifying the relevant goal makes it possible to evaluate potential measures as ones that can be plausibly argued to have an effect on it. Whether or not that is the case can then be subject to empirical investigation.

In contrast, consider the same aspect of empowerment, expanding the capabilities of the poor, but now suppose that it is viewed as an end in itself. Such a view is suggested in the *same* definition of empowerment, which also refers to the expansion of capabilities of poor people to "*participate* in . . . institutions that affect their lives" (Narayan 2002, 14, emphasis added). One could presumably argue that such participation has direct value—that it is itself

a component of the welfare of the poor. As before, a theoretical framework, such as that illustrated in function *f(.)* in equation (1), would provide the basis for constructing particular measures. Using the same example of public projects, one could postulate that increased capabilities are directly provided in the form of activities that allow the poor greater participation. One such measure of participation, and hence of empowerment, in a development project could be whether the poor take part in all the planning meetings. Note that this measure is different from the one we came up with when viewing empowerment as a means to an end. Specifically, we no longer require that this measure result in poor people "influencing" the project decisions. If one assumes that attendance of poor people at meetings is important because we care about participation in its own right, it does not matter whether the participation leads to real influence. Nevertheless, in this case it is still important to establish that the particular measure used can be plausibly argued and subsequently shown to be an indicator of participation in the project decisions.

The above examples highlight the different consequences for theories and measures of viewing empowerment as a means to an end or as an end in itself. They also suggest a trade-off between these two approaches. When empowerment is regarded as a means to achieve an outcome, the most difficult tasks are likely to be laying out a theoretical framework showing how empowerment affects the outcome that is desired (such as changes in wealth, income, or other socioeconomic indicators) and then empirically showing that the measures chosen indeed causally affect this outcome. When empowerment is viewed as an end in itself, the theory and measures are often not that difficult to develop. In the example above, the "participation in planning meetings" measure is plausibly connected to empowerment understood as increased capabilities of the poor to participate in institutions that affect their lives. In this case, however, what is more difficult to establish is that the aspect of empowerment identified has *direct* welfare value. The more specific the measure chosen—such as participation in a set of meetings—the harder it is to establish that this measure is broad enough to have direct value, or alternatively, that it *causally* affects a broader notion of empowerment that is plausibly a part of agent welfare.

The point here is not to suggest that empowerment should only be viewed in one way—either as an end or as a means to an end. In fact, it is likely that both interpretations are correct. However, it is important to make the distinction and to be explicit about which is being used in a particular study, since they imply distinct theoretical frameworks and measures and require addressing different sets of issues. In this chapter, the primary focus is on viewing empowerment as a means to an end and illustrating how theory can be used to construct potential measures and defend them empirically. While much of what follows could also be applied if empowerment were viewed as an end in itself, economic methodology has fewer tools to establish whether empowerment has direct welfare value. That alternative is accordingly less developed in the following discussion.

From Theoretical Framework to Empirical Measures

In order to establish causality from empowerment measures to an outcome of direct value—whether an aspect of empowerment itself or a socioeconomic welfare outcome—one has to start with a theoretical framework for such effects, that is, what factors to include in equations (1) and (2) and what form the functions *f(.)* and *g(.)* take. Such a theoretical framework both lays out the hypothesized causal relationships and suggests measures to construct. Once such a framework has been developed and related measures constructed, we can subject them to empirical investigation. If we can establish that causal relationships do exist, the measures can appropriately be used as policy instruments.

The best way of illustrating this is to offer a potential theory of empowerment and then show how it translates into measures of empowerment for an actual case.

If empowerment is viewed as an end in itself, then the theory should start by defining empowerment concretely and in a manner that justifies why it can be viewed as an end. The more precise this definition, the easier it will be to construct measures that capture empowerment. For example, one could claim that (part of) empowerment is the granting of "voice" to the empowered agent and that voice is an inalienable right included in an individual's overall welfare. The theory would then explain what it means for an agent to have voice. This would require suggesting both measures of empowerment in terms of voice, and factors that, while not directly measures of voice, influence voice.

If empowerment is viewed as a means to an end, on the other hand, then this theory must start by indicating the particular end of interest and then describe the processes through which empowerment affects this end. This is the process followed in this chapter, which lays out a specific theory that views empowerment as a means to achieve increased economic welfare of the agent. This theory is then applied to data collected in the Baltistan study, showing how to generate measures of empowerment and empirically test them.

The goal is not to convince the reader of the particular theory of empowerment presented, or to present such a theory in a comprehensive way, but rather to illustrate a general process that may be used for other theories as well. In other words, the intent is to show how constructing measures of empowerment should begin with a theory of empowerment, which then provides the basis for development of measures.

Theory Development

The empowerment sourcebook identifies four key elements of empowerment: (a) access to information, (b) inclusion/participation, (c) accountability, and (d) local organizational capacity (Narayan 2002, 18). These elements can be used to develop a theory of empowerment that explains how empowerment of the agent (an individual, community, group, etc.) brings about desirable

outcomes such as an increase in the agent's economic welfare or, more specifically, provision of basic services, access to justice, and improved governance.

The exercise starts by hypothesizing that the two underlying theoretical components in empowerment are information and influence, and then seeks to formalize these two concepts. Each is considered in only the most simple form, with two types of actors: an agent from the beneficiary community and an external agent. Further elaboration of such a model with inclusion of additional factors, such as who controls information and how its flow can be constrained by power brokers, falls outside the scope of this illustrative exposition.

Information

Information as a component of empowerment is conceptualized as both *provision of* information and *access to* information by the empowered agent, to and from the external agent or organization respectively. When poor people or communities are empowered, they are both able to provide information about their own preferences and gain information from outside that may in turn enhance their capacity to make optimal choices. Both types of information are likely to lead to increased welfare of the empowered agent.

Greater *provision of* information is expected to benefit the agent, as the final outcome is more likely to match the agent's needs. For example, when a decision must be made about which public project to build in a community, if an empowered community can express its preferences to the local government and make sure that the project is chosen accordingly, the project is more likely to succeed. The development literature abounds with instances of failed projects built without any local consultation (Tendler 1997), such as drinking water schemes that failed because people preferred walking to the local well (the value of time saved by an in-house tap being outweighed by the value of having a regular social space away from home).

Access to information, the result of information flows from external institutions to the agent, can also help by allowing the agent to make more informed decisions. For example, in the Baltistan study, it was surprising how often a community asked an external agency to construct a particular type of project based on inadequate information. Because communities incorrectly assumed that the external agency only provided certain types of projects and that asking for anything else would result in their not receiving a project at all, they often excluded projects that reflected their most pressing needs. Such errors might have been avoided had the community been empowered to obtain information at will from the external agency.[1]

In addition to being a separate component, information is also an aspect of the other three key elements of empowerment listed above. Participation can be partly thought of as a means of providing and gaining information. Similarly, such transfer of information is essential to social accountability and also helps foster local organizational capacity. It is important to note that the concept of information as used here is broader then simply the pro forma act of "asking and telling." Information exchange as a component of empowerment also implies that both parties are willing to supply information that is

relevant to achieving goals. Unless the agent or community is empowered in other respects as well, there is no assurance that such information exchange can effectively take place. Unempowered agents may be unwilling or unable to express their preferences fully—unwilling because they may correctly perceive that there is little chance of these preferences being met, or unable because they lack enough information to choose the best option without more direct participation in the decision-making process.

Influence

While information, as formalized above, is necessary, it is by no means sufficient to produce the desired outcome. The second component, influence or "bargaining power," is also required. Agents may be able to share information perfectly. But unless they have the ability to influence the decision and, moreover, *know* that they have this ability, they will have little incentive to either provide or gain the requisite information. Even if they do so, they have no assurance that this information exchange will actually affect how the decision is made.

Formally, this chapter defines the influence component of empowerment as the agent's "relative ownership" of a particular decision. This is based on extending the property rights concept in economics, which defines ownership of a physical asset in terms of "residual control rights" over the asset (Grossman and Hart 1986; Hart and Moore 1990), to also include less tangible "assets" such as decisions. This extension to the property rights theory suggests that ownership of such assets should be given to agents whose effort or investment is more important in influencing the return for that particular decision. The idea behind this is intuitive: by giving greater influence in a decision to the agent whose investment matters most for the decision, we are ensuring that this agent will have a high incentive to make the investment. This leads to higher overall benefits from that decision for all parties involved.

A detailed model was developed to examine the role of community participation in decision making in the context of the Baltistan study (Khwaja 2004). In the study, a local community's influence in externally initiated public projects is measured by the community's nominal participation in decisions that affect the sustainability of the project and hence the benefits accruing from it. The community's participation in a project decision is viewed as a means of empowerment to the extent that such participation brings a greater likelihood of influencing the outcome of the decision. The property rights theory can be used to evaluate which of the two agents involved, the community or the external organization, should be more empowered to make the decision. Specifically, this theory suggests that this choice should depend on the nature of the decision. For a particular decision, the agent with the more important investment in the outcome should have greater influence.

It is hypothesized that a community's influence, as measured by its participation in a decision, may be desirable in some cases but not in others. This idea stands in contrast to much of the literature, which often views community participation as an unqualified good (Narayan 1995; Isham, Narayan, and

Pritchett 1995; World Bank 1996). Deciding usage rules for a community project is an example of a decision that the community is best suited to make. Hence, empowering the community by increasing its participation would be preferred. The data in the Baltistan study confirm that greater community participation in this decision is indeed associated with better project performance. Regression analysis shows that a 10 percentage points increase in community participation in nontechnical decisions (see table 12.1) is associated with a 5.5 percentage points increase in quality of project upkeep (measured on a 0–100 scale). This result is not surprising, as there are numerous examples of development projects that failed because external agencies ignored community preferences in favor of standard blueprints or other externally imposed rules.

The theoretical framework, however, also suggests that in some cases external agencies may be better placed to decide than communities. The implication is that decisions in such instances should be *less* influenced by the local agents. An example is deciding the appropriate scale of a project, such as its physical dimensions or planned productive capacity. Such decisions may require engineering knowledge; in the context of the Baltistan study, this would imply greater investment on the part of the external organization. This prediction is indeed confirmed by the data: regression analysis shows that a 10 percentage points increase in community participation in *technical* decisions (see table 12.1) is associated with a 3.8 percentage points decrease in quality of project upkeep (Khwaja 2001, 2004). This result is robust, even when community- and project-level controls are taken into account. Thus, as predicted by the theoretical framework, empowering the community by increasing its participation in these decisions is actually associated with *lower* project performance, presumably because this community role lessens the influence of the external agency whose judgment is needed for technical decisions.

This theory and these examples are not presented as generally applicable. Rather, they are intended to illustrate how defining a theoretical framework with hypothesized causal connections facilitates the construction of appropriate measures.[2] The theory provides a structure that restricts the search for potential measures of empowerment to those that capture the information and/or influence aspects of empowerment. Moreover, while the general theory may apply to cases beyond the particular study, appropriate measures may vary from one context to another. A variable that is a useful measure of empowerment in one environment, where it captures the information an agent provides, may not be most appropriate in another environment or at a different time in the same environment, where the nature of the information transfer is different. The benefit of laying out a theory is that it provides general rules yet allows for development of context-sensitive measures, as illustrated in the following section.

Potential Measures

Taking the summary theory presented above as a starting point, the next logical step is to search for and explore causal relations between potential

Table 12.1 Participation Levels in Project Actions and Decisions: Summary Statistics

Action/decision	*Observations (N)*	*Mean (%)*	*Standard deviation*
Nontechnical			
Selecting project	132	80	29
Deciding level and distribution of community labor contribution in project construction	132	36	33
Deciding level and distribution of community nonlabor (cash) contribution in project construction	132	24	30
Deciding wage to be paid for community labor used in project construction	132	36	35
Deciding on any compensation paid for nonlabor community resources used in project construction (e.g., land given up)	119	13	25
Labor work for project construction	132	85	24
Monetary contribution for project construction	132	36	41
Deciding project usage/access rules (e.g., who gets to use the project when)	132	13	23
Deciding sanction measures for project misuse (e.g., amount and nature of fines levied)	132	14	21
Raising internal (to community) funds for project construction and maintenance	132	9	19
Deciding on distribution of project benefits (e.g., allocation of water, electricity across households)	129	19	32
Deciding on maintenance system, policies, and rules	132	20	29
Deciding on level and distribution of community monetary contribution in project maintenance	132	17	28
Deciding on level and distribution of community labor work toward project maintenance	132	28	34
Deciding on nature, level, and extent of any sanctions imposed for not participating in project maintenance	132	22	29
Overall participation in nontechnical decisions	132	30	19
Technical			
Deciding project site	132	23	31
Deciding project scale (length, capacity)	132	18	27
Deciding design of project	132	11	21
Deciding time frame for project construction	132	10	19
Raising external (to community) funds for project construction and maintenance	132	22	34
Overall participation in technical decisions	132	17	18

Note: Participation is measured in percentage terms. Thus 80 percent mean participation in the first decision (selecting the project) implies that of all respondents surveyed in all 132 projects, 80 percent said they had directly participated in the decision. In contrast, only 23 percent of the respondents had directly participated in deciding the site of the project.

measures and desirable outcomes. The following examples, in the areas of information and influence, illustrate the point.

First consider measures related to information, in the context of the Baltistan study. Suppose that knowledge available about the society and fieldwork both suggest that the traditional and effective way to elicit information from the villagers is to hold a public meeting where everyone can express his or her views openly. The implication is that a possible measure of empowerment is the extent to which all villagers have access to such public meetings. Similarly, one may judge that it is better to hold such a meeting at a traditional public gathering point rather than in an individual's home. Once such a measure has been implemented in the context of a particular project, it becomes possible to test whether such an information exchange actually has the expected effect on project outcomes.

A different context might imply a different measure of information exchange. In a more hierarchical setting, for example, information might be harder to elicit through public meetings. In such a case, the relevant question to ask is whether each villager is able to express his or her opinions individually and privately. This is better done by speaking with villagers individually in their homes rather than in a central place open to public observation. In such a context empowerment would be measured not by access to public meetings but rather by the degree to which individuals had opportunities to offer their opinions privately. Yet a third context might require setting up local coordination mechanisms so that each person is assured that others will also be willing to report. This might be the case if no one wants to "snitch" on a fellow community member, even if that member is taking undue advantage or misusing a public good.

Before suggesting potential measures for influence, it is necessary to detail the elements needed to define influence in a particular context. Influence takes its meaning in reference to a particular activity or event: one talks about influence *in* a society, electoral campaign, or project. In the Baltistan study the issue is influence in public projects.

In turn, one can subdivide such an activity into attributes, and then consider influence with respect to each one. Relevant attributes for an irrigation channel project might include the project design and length, duration and method of construction, maintenance system, and so on. Table 12.1 lists various actions and decisions that were considered in the Baltistan study as important attributes of a public project.

Once an activity is characterized by such a set of attributes, an individual's influence on a particular attribute can be defined as his or her control right over this attribute. The following example illustrates how such a consideration can help refine the measure. A meeting is held to decide which project will be built in a village. All the villagers show up at the meeting and a decision is reached after some negotiations. This may be considered empowering since all the villagers attended the meeting. Bringing rights into the picture allows a more precise perspective. One may consider both (a) attending the meeting and (b) voting in the meeting. While all villagers had the right to attend, and did in fact show up, there is no reason to presume that all had the right to vote on decisions. If the decision was in fact made by the village head, then the influence

of other villagers might have been quite low. What measure to use to capture this control right will thus depend on the context. In the context of the Baltistan study, asking whether an individual participated in a certain decision was considered an effective way to measure such rights, because the society was not too hierarchical and therefore nominal participation was believed to result in real influence.

Finally, influence measured for individual units (persons, groups, or countries) can subsequently be aggregated, with a group's influence on an activity considered to be a function of the individual members' influence. In the Baltistan study, the procedure was to conduct individual interviews and ask community members whether they or their household members had participated (either directly or through a proxy) in various decisions about the project. Empowerment at the community level was then estimated by averaging these responses. Table 12.1 gives the average participation levels for each of these decisions in all the projects surveyed. These participation levels then serve as measures of influence and hence empowerment of the community.[3]

In sum, the process proceeds in a series of steps. First one identifies the relevant attributes that define an activity. Table 12.1, for example, lists some of the decisions that affect a public project, ranging from the selection and conceptualization of the project to its implementation, usage, and maintenance. For each attribute, one then considers the control rights over the attribute and how individuals have access to those rights. In the Baltistan study, this was done by aggregating the nominal participation of a representative group of project beneficiaries in the community. Such a simple technique illustrates the approach, and can be elaborated as desired. For example, one could also introduce individuals' relative power in the group, indicated by wealth or socioeconomic status (see note 2), or consider the role actors outside or inside the group play in affecting the relative influence of group members. In this regard, mixed-methods approaches, relying on techniques other than survey questions, can be extremely useful in shedding additional light on information flows and relative influence structures.

The main benefit of the approach suggested here is to provide a structure for measuring influence without making the concept overly rigid. The researcher or practitioner must ask what influence means, whose influence is referred to, and how an activity is best described. In order to operationalize control rights, it is also necessary to consider the underlying power structure in the particular context. The approach allows flexibility. It does not presume, for example, that increasing an agent's influence is necessarily beneficial to the agent's welfare. What control rights are considered optimal depends on which agent has the most to offer to the outcome, although such judgments may be particularly difficult if the resulting reallocation of rights leads to undesirable equity outcomes.

While the discussion above examines empowerment as a means rather than an end, similar considerations of constructing measures according to a theoretical framework and specific context would apply if empowerment were viewed as an end. In either case, however, the next step of subjecting measures to empirical tests is particularly challenging. For example, if we view empowerment

as a means to an end, then we need to establish that the empowerment measures are indeed causally related to the end. Moreover, even if empowerment is viewed as an end, it is likely that the measures that one arrives at are not obviously indicative of empowerment. It may then be necessary to strengthen the conceptual case by establishing a causal link between such measures and other acknowledged indicators of empowerment. The next section of this chapter addresses potential issues in establishing such causal links.

Establishing Causality

While coming up with the appropriate theoretically based instruments is a crucial step toward formulating an empowerment-based development framework, there is another essential step. That is to *empirically* establish causality from a measure to the desired outcome (be it empowerment or a socioeconomic end). It is not possible to infer the direction of causation simply from correlations, since the measure being taken as a cause may itself be influenced by other factors affecting the outcome or by the outcome itself. Such "endogeneity" problems must be addressed in order to confirm that the observed relationship between the proposed measure and the desired outcome is in fact a causal one.

The estimate of the relationship between the measure of interest and the outcome may be incorrect because the measure is correlated to the part of the outcome that remains unexplained (that is, the measure is endogenous). There are a variety of reasons why such endogeneity problems may arise, as illustrated in figure 12.2. The figure shows both the causal relationship from empowerment to welfare factor that we are trying to establish and problems that may make establishing such a relationship difficult.

One classic problem, illustrated in figure 12.2 in the bottom-right box, is that of an "omitted" variable. This is a variable that affects both the outcome

Figure 12.2 Empowerment as a Means: The Problem of Causality

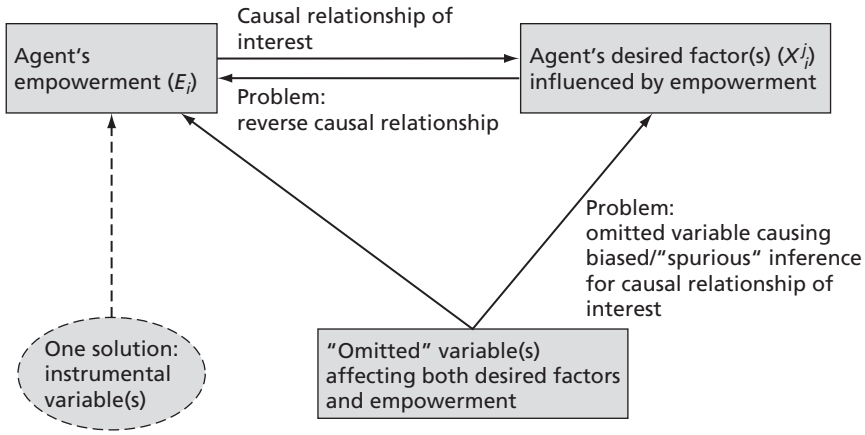

(the factor affecting the agent's welfare) and the empowerment measure. If this is not taken into account, one may erroneously assume a causal relationship between empowerment and the outcome when in fact the correlation is instead explained by a third factor—the omitted variable. In the context of the Baltistan study, suppose participation in project decisions is taken as a measure of empowerment. Our outcome of interest is project performance since it is expected that better-performing projects raise agent welfare. However, even if there is a positive relation between participation and the project performance, it may be incorrect to assume a causal relationship. Both project participation and performance may be influenced by a third factor, such as the level of organization of the community. In this case, the effect of the level of organization of the community might be mistakenly identified as the effect of participation. In an extreme case, participation may have no effect on performance at all; the relationship may be entirely spurious, that is, entirely due to the factor(s) not specified in the model. While this problem can be corrected by simply including this omitted factor, such variables are often hard or even impossible to measure even when they are identified.

Alternatively, as also illustrated in figure 12.2, the actual relationship may be fully or partly the reverse of what we imagine, that is, the empowerment measure may itself be affected by the outcome of interest. Continuing with the earlier example, it could be that if a project is doing well, then participation in the project increases. In this case, the causal relationship is not that participation of the community *causes* a project to perform well but rather that a well-performing project *attracts* community members and increases their participation. This problem can partly be addressed by collecting information on participation in decisions that took place prior to a project's performance being revealed. However, as such information is often collected by using recall data, one may still get reverse causation through "halo" effects; that is, a respondent may tend to report greater prior participation for projects that are currently doing well. One possible solution, used in the Baltistan study, was to show that such halo effects are unlikely by demonstrating that respondents who declared higher project participation were not more likely to report that the project was in a better state. Such solutions may not be possible in many cases, and instead more general empirical techniques will be required.

There is a wide array of empirical methods intended to address such concerns, ranging from careful experimental design to the use of instrumental variables and fixed-effect estimations. Each tries to solve the above problems by establishing or creating exogeneity of the causal factor of interest. In other words, one tries to construct situations where it is clear that the factor in question is not correlated with the unexplained part of the outcome—that is, it is "exogenous."

In randomized or experimental design techniques (prevalent in the medical sciences) one creates such exogeneity by randomly assigning a "treatment" (the factor of interest) to some groups but not others. In the Baltistan study this would have involved selecting some projects at random and introducing community participation in them but not in others. However, carrying out such experimental studies is often not feasible for practical reasons or is undesirable on ethical or fairness grounds.

The instrumental variables approach relies on identifying a measure or instrument associated with the causal factor of interest—the measure of empowerment—that is known or assumed not to be causally linked to the outcome. For example, for the participation measure, it could be the case that among two projects implemented by the same external agency, one took place before a change in the agency's policies in the direction of greater participation. This change in the agency could then be used as an instrumental variable for participation. However, one also has to argue that this instrument only affects the outcome through the particular measure of empowerment and not directly or through any other unobserved factor. These conditions are relatively stringent, and it is often difficult to find suitable instruments.

Finally, a "fixed effects" approach tries to get around these problems by forcing a comparison *only* across projects for which there is unlikely to be an endogeneity problem. This approach was used in the Baltistan study. Consider again the relationship between participation (the empowerment measure) and project performance. To respond to the concern that more "organized" communities have both higher participation and better performance, one could only compare projects within the same community. Since the omitted community-level variable would be the same for all projects in a community, it could not be responsible for any observed differences. Such an approach, however, would of course not compensate for possible omitted variables at the project level.

The point here is not to recommend these particular techniques or others that could be added. It is rather to stress that causality is an essential issue to address in evaluating measures of empowerment. A plausible relationship between an empowerment measure and a particular outcome is not adequately verified by observing a correlation between the two. Such a correlation is simply an *association* and *not necessarily a causal relation.* One can advance toward establishing causality by considering other possible reasons that the measure of empowerment may be related to the outcome. Often some concerns can be allayed by a careful examination of the context or by drawing on other quantitative or qualitative sources. Remaining concerns may be addressed, to whatever extent possible, by use of the empirical techniques outlined above. The degree to which all concerns can be addressed will, of course, vary. However, any researcher or practitioner must at least acknowledge and discuss these concerns and evaluate how serious are those that are not fully addressed.

Conclusions

In sum, in order to develop an empowerment framework for development that can be implemented, it is necessary first to distinguish between aspects of empowerment that are considered of direct value, that is, as ends in themselves, and those that are means to an end. If one takes the first view, then there needs to be a justification that a particular aspect of empowerment is itself valuable, in sufficient detail to allow for specifying measures of empowerment and then empirically establishing that they indeed causally affect empowerment.

If empowerment is instead viewed as a means to specific ends, such as commonly valued socioeconomic outcomes, then a theory is required specifying how empowerment is hypothesized to affect these outcomes, followed by procedures to establish empirically that the empowerment measures are causally related to such outcomes. In this chapter, the second alternative was illustrated by developing an outline of a theory of empowerment in terms of information and influence affecting project performance outcomes. The theory was then combined with context-dependent knowledge to develop measures.

The chapter made use of an economist's perspective and focused on survey-based measures in particular. However, the framework developed should also prove amenable to the use of other methods. The critical point is that any measures developed, whether survey-based or otherwise, must not only be based on theory and specified plausibly, but must also be tested by methods that attempt to establish causality. Such a process, moving from theory to measures to empirical tests, may seem overly exacting. But it is essential if empowerment is to avoid the fate of previous development slogans that have been misunderstood, misapplied, and eventually discarded.

Figure 12.1 illustrates the two possible relationships between empowerment and the desired outcome, increased welfare of the agent. These correspond to equations (1) and (2) mentioned above. In relationship (1), empowerment is directly included in agent welfare, along with other factors. Since this relationship is definitionally true—empowerment is defined to be part of agent welfare—it is represented by a circle. However, even in this formulation, whether a particular *measure* affects empowerment is unlikely to be definitionally true and would need to be causally established (this can be illustrated by a directed arrow from the measure to agent empowerment but has not been included in the figure to keep it simple). Relationship (2) shows empowerment (or any of its measures) as affecting factors that are components of agent welfare. Since this is a hypothesized causal relationship, it is indicated by a directed arrow.

Figure 12.2 highlights the problem of causality where one is trying to infer that a particular measure of empowerment causes an increase in a factor that affects agent welfare (such as agent wealth). The figure illustrates the causal relationship of interest but also shows two problems, reverse causality and omitted variables, that may lead to difficulties in establishing the correlation between the empowerment measure and the welfare factor as causal. It also illustrates one possible solution of using instrumental variables (that only affect the empowerment measure but do not directly affect the welfare factor or the other omitted variables).

Notes

1. For 9 percent of all projects surveyed, community members said that they had chosen the project from a list of projects they "knew" the external agency provided or had suggested. This proportion is quite large given that the external agencies we

examined were, if anything, far more careful than is typical about using participatory tools to elicit community preferences. Interestingly, another 10 percent of projects were chosen either because they were perceived as having been suggested by the external agency or because they had been chosen by a neighboring village.

2. Community heterogeneity has a detrimental effect on overall project upkeep while inequality across community members has a U-shaped relationship, with initial increases lowering project upkeep (Khwaja 2001). The theory outlined here did not explore empowerment *within* a community (i.e., empowering one community member relative to another). However, one can envisage a framework in which community heterogeneity and inequality would affect a member's influence and hence degree of empowerment *within* the community. This may in turn affect how well a collective task, such as upkeep of a public project, is performed.

3. In fact, the participation measures can also be thought of as measures that capture the informational aspect of empowerment, since it is likely that greater participation of the community also means that it is providing and receiving more information. These measures are divided into those relating to nontechnical decisions and those relating to technical decisions. This is because the theory suggests that decisions that benefit more from local information (nontechnical decisions) are best decided primarily by the community, while decisions requiring technical inputs are best left to those who can offer such technical advice. In the context of this study, the latter implied participation by an external agency's engineers. The study confirmed the validity of making such a distinction: while community participation in nontechnical decisions increased sustainability of the local public good, community participation in technical decisions actually decreased sustainability. Without the theory, we would have aggregated community participation in all types of decisions, thereby missing the two opposite effects. Thus the theory was essential not only in arriving at each measure, but also in telling us how they should be grouped.

References

Grossman, Sandy J., and Oliver D. Hart. 1986. "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration." *Journal of Political Economy* 94 (4): 691–719.

Hart, Oliver, and John Moore. 1990. "Property Rights and the Nature of the Firm." *Journal of Political Economy* 98 (6): 1119–58.

Isham, Jonathan, Deepa Narayan, and Lant Pritchett. 1995. "Does Participation Improve Performance? Establishing Causality with Subjective Data." *World Bank Economic Review* 9 (2): 175–200.

Khwaja, Asim Ijaz. 2001. "Can Good Projects Succeed in Bad Communities? Collective Action in the Himalayas." Faculty Research Working Paper RWP01–043, Kennedy School of Government, Harvard University, Cambridge, MA.

———. 2004. "Is Increasing Community Participation Always a Good Thing?" *Journal of the European Economic Association* 2 (2–3): 427–36.

Mansuri, Ghazala, and Vijayendra Rao. 2004. "Community-Based and -Driven Development: A Critical Review." Policy Research Working Paper 3209, World Bank, Washington, DC.

Narayan, Deepa. 1995. "The Contribution of People's Participation: Evidence from 121 Rural Water Supply Projects." ESD Occasional Paper Series 1, World Bank, Washington, DC.

———, ed. 2002. *Empowerment and Poverty Reduction: A Sourcebook*. Washington, DC: World Bank.

Tendler, Judith. 1997. *Good Government in the Tropics*. Baltimore: Johns Hopkins University Press.

Uphoff, Norman. 1996. *Learning from Gal Oya: Possibilities for Participatory Development and Post-Newtonian Social Science*. London: Intermediate Technology Publications.

World Bank. 1996. *The World Bank Participation Sourcebook*. Washington, DC: World Bank.