Whose Policy Counts? The Case for Policy That Listens

Vijayendra Rao, Northwestern University and CIFAR

In 1992, India amended its constitution to create some 250,000 elected village governments covering roughly 800 million people — the largest exercise in democratic decentralization in human history. The effects of this 73rd Amendment have varied enormously across states. In Kerala, an extraordinary People's Campaign in 1996 created a system of nested participatory planning. Street-level groups deliberate on their needs, which are passed on to a ward-level assembly, then percolate upward to a village assembly where all expressed needs are prioritized and matched with budgets. Over the course of the year, the agreed-upon projects are implemented and monitored in active collaboration with citizens.



*Neighborhood meeting (*Ward Sabha*) in Kerala, 2003 @Vijayendra Rao*

In other states, the same assemblies, created by the same amendment, are ceremonial rooms where nothing important is decided. In the 300 village-meeting transcripts that Paromita Sanyal and I analyzed for "Oral Democracy," we found that differences in literacy, caste or wealth inequality did not explain how well these village assemblies worked. They were driven by whether the state government actually wanted decentralization to work – whether it devolved real money and real authority to the assemblies. The same constitutional amendment produced very divergent realities.



Gram Sabha in Tamil Nadu, 2012 @Vijayendra Rao

But understanding this divergence in the quality of deliberation has not been prioritized by mainstream economists and policymakers — because it would be a measure of process. Over the last century, the policy options on offer from economists for tackling poverty and inequality have overwhelmingly been focused on outcomes — resulting in what I call *outcome-policies*.

An outcome-policy manipulates the levers of government — transfers, taxes, expenditures, regulations — to hit a target and is judged by whether it hits it. Conditional cash transfers are judged by whether children stay in school; estate taxes by whether they improve intergenerational mobility; health insurance by its effects on health and hardship.

Outcome-policies are tied to metrics, which is much of their appeal. When the politics are favorable, they are relatively easy to enact by legislative or executive action. They are also easy to reverse.

A *process-policy* is a different animal. It tries to change how change happens: to equalize power in how decisions get made, each increment building on the last. Its proximate effects are subtle and hard to measure, and its full impact takes years to show. Precisely for those reasons, changes from process-policies tend to last longer than those from outcome-policies. Process policies change the machinery that produces outcomes rather than the outcomes themselves. Constitutionally protected deliberative bodies, participatory budgeting systems, adaptive systems of implementation — these are process-policies.

In the essay that preceded this one, I argued that development economics research has built a wall between itself and the people it studies. Policy has the same distance problem that research does, but the stakes are higher, because policy can directly and substantially affect people's lives at a scale that research cannot. The distance, though, is not just a matter of bad implementation. It is built into the frameworks we use to think about what policy is.

The Three Dominant Ways of Thinking About Policy

Policies for welfare, growth, inequality and poverty are conditioned by three broad paradigms, all mostly derived from economics.[1] All three are outcome-oriented at heart.

The *neoliberal* paradigm starts from deep skepticism that social science can understand complex human interaction well enough to engineer a better world. So it falls back on first principles: laissez-faire markets, free trade, human rights, electoral democracy. Free markets deliver growth; elections keep governments honest. The proponents of this paradigm have some empirical evidence for this: over the long run, democracy may indeed cause growth, and there is evidence that growth is an important determinant of poverty reduction.

This is the view that underlay the Washington Consensus of the 1980s and 1990s. Its answer to poverty is to maximize growth and equalize access to human capital, and there is something to it — but the effect of growth on poverty depends on how unequal a place is to begin with, and there is strong reason to believe that neoliberal policies increase inequality and erode social support. This can be by, for instance, creating opportunity markets that commodify access to basic needs like education and housing. Also, note that the largest reduction in poverty happened in China, which is not a democracy and grew and reduced poverty by making policy adaptively and experimentally, adjusting as it went.

The deeper limitation of neoliberal policymaking is disciplinary. Rational-choice individualism has no comfortable place for social exclusion or cultural and political inequality, so every socially driven argument — for gender equality, for social cohesion — has to be smuggled in through a growth prism. (Hence the slew of 1990s papers arguing

[1] These are "ideal types" and may blur into each other in practice.

that women's education is good for GDP.) Welfare gets measured in income and consumption and poverty becomes a line denominated in dollars and cents, and a large survey industry exists to determine it.

The *neo-Keynesian* paradigm is the counter. Markets fail — through externalities, information asymmetries, market power — and the social sciences can tell the state how to correct them. Growth still matters, but there is acute awareness that markets do not deliver basic services to the poor and that discrimination is systemic. Household surveys now collect data on gender, religion, race and caste; equality of opportunity is a stated goal, backed by law and affirmative action.

The result is the intellectual justification for interventionist government — the London Consensus is the most recent comprehensive statement of this, with its slate of minimum wages, progressive property taxes and public investment in everything from infrastructure to nutrition. Along the way, the paradigm has developed a deep knowledge of government failure and clientelism. But even as it has sought to understand government failure, it does not seek to fix it through process. Instead, it builds on expert diagnoses by economists and social scientists — the state corrects, the citizen receives.

The third paradigm — *neopaternalism*[2] — is arguably a more stringent extension of the second, and it has exploded in influence over the last two decades. It starts from the opposite end from neoliberalism: an abiding faith in the power of social science and data — behavioral experiments, randomized trials — to give direction to policy.

It takes its cues from public health, with a firm commitment to "evidence-based" everything. Its behavioral foundation is the research on scarcity: the poor face such punishing constraints on time and attention that they are steered into suboptimal choices, so the freedom to choose means little when the ability to choose is itself constrained. That is the justification for the top-down push — the nudge — away from actions, norms and practices that the policymaker has determined are harmful.

I want to be fair to this world, because it is the one I have been surrounded by and worked in for decades, and its advantages are real. The trials generate vast amounts of data about the conditions of the poor, well beyond income and consumption. They create a culture in which policy has to be justified with evidence rather than hope or idealism. And they weed out the bad intervention ideas.

But the disadvantages of this model of the world have been well catalogued, and two matter most here.

The first is overreach: policy driven by the demands of scientific technique rather than technique driven by the needs of policy — the tail wagging the dog. If all evidence must pass through a randomized trial, interventions that cannot be randomized disappear from the menu. And because a good trial needs good implementation — a messy rollout contaminates the design — trials effectively assume away what Albert Hirschman argued

[2] To quote Thaler and Sunstein, "a policy counts as 'paternalistic' if it is selected with the goal of influencing the choices of affected parties in a way that will make those parties better off."

60 years ago is the central problem of development policy: the complicated process by which ideas become action on the ground.

The second is context. An intervention that works in one place routinely fails in another. The randomizers' answer is to run trials across many settings and at much larger scale, which works for simple interventions but breaks down for complex ones — the interventions tailored for, and especially by, communities, where designs vary too much for any sample to have adequate statistical power.

So neopaternalism drifts, by the logic of its own tools, toward interventions that are simple, uniform and predictable. It is a paradigm far more comfortable with certainty than with reflexivity. Culture enters the analysis as a set of "traits" to be scored good or bad; the study of social life becomes part of the expert's apparatus for nudging people toward welfare as the expert defines it. Nobody asks the poor what they want. And it is not a rhetorical flourish to say that. It is a fairly precise description of the method.

As I write, British and American aid budgets have been gutted, and two economists I respect — Rachel Glennerster and Dean Karlan, who served as chief economists of the two countries' aid agencies — have jointly responded with a call for "radical simplification": take what money is left and spend it on a short list of proven, cost-effective "smart buys" — malaria nets, nutrition supplements, structured lesson plans, cash transfers — at large scale, shorn of the complexity that makes programs "expensive to design, hard to manage, and almost impossible to evaluate." They have a point, up to a point.

Mission creep is real, and it is often driven by donors, as anyone who has watched a project buckle under the weight of a phalanx of consultants knows. Reducing mission creep and focusing on what you meant to do in the first place has value. But look at the third item on their charge sheet. "Almost impossible to evaluate" is a statement about the evaluator's preferences and their expert-driven worldview; it prompts the question of "evaluable by whom?"

When the menu of interventions is restricted to what a randomized trial can certify, it fills up with simple, uniform, predictable interventions, and the complex policies with subtle effects and unpredictable trajectories — which can displace the steady state and shift to a better development path — gradually fall off it. But many effective programs are not simple and they can take decades to bear fruit. Indeed, it would be impossible to price the net gain per unit.

But consider who paid for these programs. The policy reforms of Kerala were not financed out of the bilateral aid budgets (now being gutted); they were political projects of Indian state governments, built (largely) with domestic money and domestic legitimacy. The principal financier of process-policy is (often) the national state, not the donor.

Radical simplification may be a good way to triage for shrunken bilateral aid budgets, but it can say much less about what governments and multilateral institutions should do with their own budgets.

Seeing Like a Policymaker

Why do three paradigms that disagree about almost everything else converge on focusing on measurable outcomes?

James Scott's "Seeing Like a State" showed why: states can only govern what they can "see," so they work hard to make society "legible" — standardized categories, uniform measures, administrative grids.

There is nothing inherently wrong with this. Without legibility, there is no taxation, no public health, no welfare state. But the categories are not just descriptions; they are instruments that remake the world in their own image. And what gets stripped out along the way is what Scott, borrowing from the Greeks, called *métis* (common sense). This is practical, local, experience-based knowledge, provided by those who know the area and the place.

His examples have become famous. German foresters replaced messy, diverse woodlands with neat rows of Norway spruce that could be counted and taxed. Brasília's grid-based urban structure, shaped like an airplane, replaced the messy reality of settlements that grow organically.

Legibility is all over the policy world. A market price is legibility: it compresses a tangle of preferences and systems of production into a single number. A poverty line is a legible expression of poverty — which is a complex, multidimensional experience. Randomized control trials are also an exercise in legibility — making behavior legible by stripping away context and subjecting lives to a treatment variable.

None of the above paradigms disagree about legibility. They argue only about who should make the policy — the market, the ministry, the experimenter.

And none of them ask the people, who are ostensibly the beneficiaries of the policy. The outcomes may be legible, but there is little effort to make them legible to those the policy affects. Seen from this angle, radical simplification is not radical at all — it is the neopaternalist paradigm's default option made legible by a budget crisis. In some ways it represents a return to the expert-driven, top-down vision for policy that organizations like the World Bank were criticized for pushing on the world in the 1970s and 1980s.

Elections Are Not Enough

All three paradigms have the same problem — who checks the policymakers? All three paradigms also have the same answer: democratic electoral accountability.

Neoliberals and neo-Keynesians lean on universal franchise to keep governments aligned with citizens' interests. Neopaternalism leans harder still: Cass Sunstein's libertarian paternalism is explicitly premised on democratic constraints keeping the nudgers benign.

But the electoral check is not perfect. Elite capture and clientelism are pervasive. Democratic backsliding is relatively common. And we now live in an age of surveillance

capitalism, in which the capacity to track people with extremely precise data threatens to delegitimize democratic processes altogether. If the ballot box is the only check the three paradigms recognize, then as the ballot box weakens, so does the legitimacy of the entire policy edifice built on top of it.



Jeep used for a village council election campaign, Karnataka, 2003 @Vijayendra Rao

This is not a problem confined to poor countries. The question of who checks the policymakers is becoming urgent everywhere — from the United States to Hungary.

Meanwhile, civil society is not limited to the ballot box. Across the world, at every level, people have organized themselves into bodies that monitor power between elections — what John Keane calls monitory democracy: citizen committees, self-help groups, minipublics, social audits, watchdog NGOs. Such forums sometimes emerge organically, but they can also be created: development organizations have spent hundreds of billions of dollars inducing participation.[3]

These institutions exist even in authoritarian settings. In rural China, citizens' groups working with sympathetic civil servants have built accountability without democracy inside

[3] With admittedly mixed results.

systems of authoritarian deliberation. In other words: a paradigm of action that runs on process rather than outcomes already exists. It is time to give it a name and a description.

Reflexive Policy

I call the fourth paradigm *reflexive policymaking*. Reflexive is used in the sense the word carries in sociology and anthropology: a reflexive practice is one in which the professional treats the boundary between herself and the people affected as a problem to be worked on, not a wall to be maintained.

The idea is a close cousin of Robert Goodin's reflective democracy: an effective democracy has to cultivate its citizens' capacity to "deliberate within," to run other people's situations through their own heads. Deliberating within is what makes cognitive empathy possible at all. The policy challenge is to produce it at scale: to redesign decision-making so that politicians, officials and citizens develop cognitive empathy for one another, horizontally and vertically.

Grand as that sounds, the practice is concrete, and pieces of it are operating all over the world. Four processes recur.

The first fills in the blank space between elections with continuous dialogue between citizens and governments — deliberative systems anchored in some officially sanctioned civic space, a forum or a regularly scheduled public meeting, where an ordinary citizen can directly influence the decisions that bear on her life. Forums alone are not enough; they have to sit inside a culture of dialogue in which speaking *and listening* become everyday practice. That is what Kerala built, and its neighbors did not.

The second equalizes power in decision-making directly, by giving voice to disadvantaged groups — and, more radically, by reserving seats for them, as India has done for women and lower castes in its village governments. Reservation is itself an admission that deliberation carries no automatic guarantee of equity — an all-male council can ignore women's concerns, and a government run by upper castes may ignore the concerns of lower castes, no matter how many times they are brought before them.[4]

Process-based policy does not assume that people will always govern justly; rather, it claims that for a government to be legitimate, it must be more directly run by the governed. They are better placed than the experts to define the problems about which they have more information and are better placed to address them; having citizens determine foreign policy is obviously not optimal, but having them run local schools is.[5]

---

[4] There is a very large (and growing) literature on women's and caste reservations in India that show that they are — on balance — equalizing both on policy choices and deliberation. This paper summarizes the literature.
[5] Decision-making should, as much as is within the realm of feasibility, follow the principle of subsidiarity: "central authorities should perform only those tasks which cannot be performed effectively at a more intermediate or local level."



Village council president, Ujjire, Karnataka, 2003 @Vijayendra Rao

The third builds feedback loops, within government and between government and citizens, so that decision-making becomes adaptive: try, learn, adjust — instead of design, implement, evaluate, stop.

The fourth uses technology to give people control over their own data, to strengthen collective action and equalize access to public services rather than to surveil. The current excitement about AI-mediated deliberation falls under this category too, with this caveat: use the machines to scale our ability to listen and learn, not substitute for it.

The Sandwich

The capacity to participate is not something that communities alone can control. It is a political relationship spread across three layers.

There has to be a responsive state at the top, with the political will to devolve real power and the fiscal capacity to back it with real money. There has to be a motivated and organized citizenry at the bottom — cohesive enough to act collectively and deliberative enough to work through disagreement. And there has to be a working middle: the district

office, the field facilitator, the state ministry — an intermediary tier that translates citizens' voice upward, and the state's authority downward, thereby protecting local citizens from having their interests captured by local elites. Ghazala Mansuri and I, borrowing from Jonathan Fox, called this the "sandwich" in "Localizing Development."

Every case of process-policy that works at scale has all three layers, and every visible failure has a hole in one of them. This is the pattern that Mansuri and I found in reviewing the evidence on participatory development, which has a large, mixed and often disappointing record. What was clear, however, was that the successes and failures could be explained — where the top was missing, participation degenerated into empty consultation, and the community meeting became a place for cheap talk. Where the bottom was thin, collective action became difficult to organize. Where the middle was broken — and in many countries the middle is the weakest layer — elite capture became much more likely and even well-designed policy harder to deliver.

This, in retrospect, provides insight into why so much donor-induced participation has disappointed. Donors build the bottom — generally the cheap layer, the one you can deliver on a project cycle. But they skip the political work at the top and the institutional work in the middle. A sandwich with one layer is not a sandwich — and the whole thing fails.

Kerala, to return to the opening puzzle, is what it looks like when a state builds the whole sandwich. The People's Campaign organized the bottom, while the ruling coalition supplied the political will and intellectual direction for the top layer. And it did one other thing: it staffed and funded the cadres in the middle. Participatory planning works in Kerala because the street-level groups feed into the ward-level assembly, which in turn informs the village assembly. Each layer of the sandwich is linked in a feedback loop with the others.

The sandwich applies everywhere. Indonesia's Kecamatan Development Program (KDP) began in 1997 as a pilot in 48 villages, having been designed by the anthropologist Scott Guggenheim and colleagues in Bappenas, the Indonesian planning ministry. It grew out of an ethnographic study of how Indonesian villages organized themselves, which found that the culture of *gotong royong* (mutual assistance) in Indonesia created village-level institutions with a much greater capacity for durable collective action than the user groups that projects kept creating.



Mosque built with gotong royong *in Kalimantan, 2005 @Vijayendra Rao*

Its design was not dissimilar to Kerala's People's Campaign. Block grants went directly from the central government to the *kecamatan*, the subdistrict. From the neighborhoods and hamlets up, villagers deliberated on proposals and sent a delegation of at least six representatives, three of them women, to an inter-village forum where they, drawing on *gotong royong* norms, deliberated and voted on which proposals to fund. Once the forum decided, nobody higher up could alter the decision. Facilitators recruited from NGOs and the private sector took the information and worked with village organizations to implement the agreed-upon plan, circumventing the village head's tendency to capture such decisions. As Guggenheim puts it, the state's job was to do less rather than more: keep the money moving, audit it and stay out of local decision-making.

That is the sandwich. Reform-minded bureaucrats in Bappenas and the Ministry of Finance at the top, with the World Bank supplying money and air cover. A tier of trained facilitators in the middle. And at the bottom, villages with long traditions of collective action and their own ways of managing and implementing a common fund. The layers of the sandwich held long enough for norms to shift. KDP reached 2,000 villages by 2002, 6,000 by 2006, when it was renamed PNPM Mandiri, and 60,000 by 2011. In 2014, Indonesia turned it into the

Village Law, which made a block grant to all 75,000 villages a permanent entitlement in the national budget. A donor-supported experiment became a permanent part of the state.

Meanwhile, Brazil's Porto Alegre is an example of the reverse. Its participatory budgeting was a celebrated case of citizens directly deciding how public money gets spent. But after the Workers' Party (PT) lost the mayor's office in 2004, the project was largely dismantled. The bottom was still there; the top lost interest, and the middle was rendered ineffective. This is why reflexive policy requires donors, organizers, politicians, bureaucrats and citizens working together long enough for norms and expectations to shift and reach a sustainable equilibrium.

A Revolution by Stealth

Consider Jeevika, the rural livelihoods program launched in 2006 in Bihar. The state was then, by almost any measure, the hardest place in India to attempt anything of the kind. Rural poverty stood at 55.7% against a national average of 41.8%. Per capita income was a third of the national average. Only 37% of women were literate, only 1% had access to microcredit and two-thirds of rural household loans came from moneylenders charging 100 to 200% a year.[6] Alleviating rural poverty in this context would be a tall order.

Furthermore, Bihar was a difficult, highly unequal, patriarchal and caste-ridden context. Into this, Jeevika introduced something deceptively modest: self-help groups (SHGs) of 10 to 15 women, meeting weekly to save and lend among themselves, subsequently federated into village organizations and cluster-level federations. Shruti Majumdar, Paromita Sanyal and I spent four years doing qualitative fieldwork (embedded with a randomized control trial that ran for 18 months) on how the project worked between 2010 and 2014. This fieldwork became our new book, "*Revolution by Stealth"*.

Jeevika offered cheap credit, insurance and income generation, with no language of women's rights or resistance anywhere in sight.[7] In a short period of 18 months, the program refinanced 28% of household debt into SHG loans that charged a much lower interest rate of 24%, undercutting the previously available interest rates by 150 percentage points or more. There were also "general equilibrium" effects — competition from Jeevika reduced the rate charged by moneylenders by 35%.

Jeevika did not present itself as a challenge to the local order. It did not name a villain. Its targets were two things: poverty and the moneylender. It explicitly brokered the acquiescence, if not the acceptance, of the village head (*mukhiya*) and the other local power brokers by offering the village a better life without creating a revolution.

---

[6] See the first chapter of *Revolution By Stealth* for citations and broader discussion on conditions in Bihar in 2006.

[7] Women's SHG programs in India are government-run non-profits that should not be confused with microcredit or microfinance programs. As we document in *Revolution by Stealth,* group-based credit and savings are just the start, used as an initial incentive to develop a federated community of women — running into the millions — creating women's social capital that is in turn used as a platform to run all manners of "vertical" interventions including nutrition, health, schooling, food security and income generation.



Women's SHG, Muzaffarpur, Bihar, 2010 @Vijayendra Rao

The gender transformation came (much) later, and it came not from persuasion or an explicit effort to change norms, but from repeated practice that turned into ritual. Jeevika brought women together, regularly, in groups with rules. Week after week, they would meet, sing, save together; they'd lend to each other, deliberate, manage their accounts, bargain. Week after week, they'd meet. Participation built confidence and women's voices. This repeated practice built radically new social roles and reworked relationships within households and communities.

As one bewildered male respondent quoted in our book put it: "(Chief Minister) Nitish (Kumar) has turned women into men."



Village moneylender in Bihar, 2012 @Shruti Majumdar

That was the sandwich. The chief minister embraced Jeevika early and gave it sustained political protection from the top. It attracted the best young Indian Administrative Service (IAS) officers, professional staff and a network of facilitators to run the middle. And the *didis* themselves — the women of the SHGs, federated upward — became the organized bottom.

Jeevika now reaches more than 14 million women in over a million SHGs across Bihar. It is part of a national program, which now serves some 86 million women across India. This is arguably the largest federation of women in the world.

There was some resistance, of course, and there still is. But over the four years we came to read the backlash differently: it is not a sign that things are failing; it is what change looks like when power and identity are being renegotiated.

We also came to think differently about culture. It is not a set of cognitive biases to be nudged away. Culture can be changed, but it is a slow process; change requires protection from local elites and political interference.

Nothing about Jeevika looks like a radically simplified intervention. Almost everything about it is process-policy — no single outcome target, no lever pulled from the capital, just a rearrangement of when, how and how often people meet, who meets whom, who holds money and who gets to speak. And it produced cultural change in gender relations. Evaluations that revisited the original RCT sample almost 10 years later found that the program had caused 126% more low-caste women to run for political office and improved women's ability to share risk, while radically transforming credit markets in one of the most patriarchal places on earth, on a timescale the standard paradigms would have said was impossible.

Doing by Learning

None of this is easy to do at scale, and much of it is not even easy to measure. That is rather the point, though. Process-policies shift trajectories in unpredictable ways, so evaluating them means watching processes move: how decisions get made, how power relations shift, how narratives and discourse change.

These things do not sit in administrative data and clean survey answers. They require qualitative work and narrative analysis coordinated with quantitative methods in genuinely mixed-method designs, which are not the norm in development economics.

For several years at the World Bank, colleagues and I tried to build this way of working into the machinery of Jeevika and its Tamil Nadu counterpart, through something we called the Social Observatory.

The idea was to move researchers from being overqualified outsiders who fly in whenever a survey needs doing to something like the project's own think tank, with the aim of improving implementation by creating a culture of doing-by-learning. A team of economists, sociologists, political scientists, public health specialists and data scientists embedded within the project helped build decision-support systems so that staff, from the village level up, could act on good, fast data. (This was instead of yet another management information system nobody paid attention to.)

Qualitative process monitors visited a rotating random sample of villages every quarter and fed what they found back into design. Impact evaluations varied by how the answer would be used — quick impact evals when the project needed an answer soon, long mixed-method ones when it needed to understand outcomes and mechanisms.

The aim was to turn research from a tool for assessment into a tool for implementation. A paper that we originally wrote in 2015, drawing on the qualitative work, was critical of how the project's second phase was being implemented, showing that the reflexive, co-produced process followed in the first phase was being neglected in the pursuit of higher enrollment numbers. The project responded with a concerted effort to give field staff more room to use their judgment. It improved the management information system and the process-monitoring system, strengthened the feedback loops between them and the facilitators, and created forums where facilitators could learn from and advise one another. The result was that facilitators moved away from an exclusive focus on the number of women enrolled toward a more adaptive and reflexive way of working with each village.

Policy, done this way, becomes constant adaptation rather than a one-shot deal. (Incidentally, this also protects against policy stasis: it kept the program responsive to changing conditions on the ground.)

What Would It Take To Do This?

Now: the honest accounting of such policymaking.

Reflexive policymaking requires very long time horizons. It is a "slow-cooked" way of doing policy, as opposed to the simple, easily scaled McDonaldization that pervades policy work. It requires a certain tolerance for mess, because reversing inequalities in power and agency is a contentious, uncertain business.

But the deepest obstacle to implementation of process-policies is that governments — and quasi-governments like the World Bank and donor agencies — are not set up for process-policies. The logic of bureaucracies, and of the political environments they answer to, pulls toward interventions that produce quick, measurable impact without threatening the status quo. Process-policies fail on all three counts: they are slow, they are hard to measure and they threaten the status quo more or less by definition.

Part of the problem may be how policy professionals are trained. (There are other equally important reasons.) Schools of public policy are generally staffed (with the occasional infusion of a "professor of practice") by academics trained in economics or another social science discipline. Their knowledge of policy comes almost entirely from bringing academic expertise to implementers who choose to seek their advice, not from being embedded in the world of policy. So the courses that policy schools teach are, for the most part, variants of academic economics or political science, subject to the same fads as the academic departments and to the same pathology of belief: that to be a good policymaker you have to be an academically gifted expert immersed in the latest techniques and theories. There are very few courses that teach you how to sit through a village meeting, and only a small minority focus on the nitty-gritty of process issues.[8]

In practice, there is already some process-policy in place. Almost all policy in any country involves some amount of all four paradigms. Markets do coordinate dispersed knowledge;

[8] Some notable exceptions are courses taught by Matt Andrews, Erica Chenoweth, Archon Fung and Marshall Ganz at the Kennedy School, and Dan Honig at Georgetown's McCourt School.

states do correct market failures; evidence can improve design. My claim is not that reflexive policymaking should replace all other types of policymaking. Rather, it is narrower: the three dominant paradigms would all benefit from more reflexivity. Reflexivity is neglected precisely because it is hard to measure, and in the current policy culture, what is hard to measure gets treated as unserious.

Process-policies mostly come into the world through years of unglamorous effort by activists working at a small scale in social movements and NGOs. Given what is happening to democratic institutions, and to inequality, it is time to bring them into the mainstream of the policy craft.

In concrete terms, every large program should budget for embedded learning. The structure of that learning should not be determined in advance; it should be process-oriented. It also does not suffice to only do randomized trials that give us a measure of outcome at the end of the project cycle.

Donors must also fund the middle tier — the facilitators, the district staff, the connective tissue — instead of mobilizing communities and walking away. Projects should be given time to evolve and adapt rather than be judged on whether they follow a predetermined trajectory of change. And schools of public policy should require — alongside the econometrics sequence — that their graduates actually spend time in the places where they will likely work. None of this is impossible, or even particularly difficult, to do. It is merely beyond current practices.

Whose policy counts? The three dominant paradigms answer, in effect: the market's, the ministry's, the expert's. The people the policy is for appear as consumers, recipients and subjects, respectively; they are not shapers of policy.

There may be a different answer. Policy counts when the people it is meant to serve can shape it, contest it and correct it — when the woman who was dragged out of my focus group in Karnataka, whom I described in my previous essay, has somewhere to take what she knows. Research that does not listen produces papers about a world it has not understood. Policy that does not listen produces programs that are neither tailored to people's needs nor able to adapt across place and time. Policy that listens can shift the political, social and cultural equilibrium — and that is what makes change last.

Vijayendra Rao is a development economist and social scientist who combines econometric methods with ethnography. He spent 27 years as a Lead Economist in the Development Economics Research Group at the World Bank and is currently the Roberta Buffett Distinguished Scholar-Practitioner in Residence at Northwestern University. His interests include gender, culture, participation, deliberative democracy, political economy and innovations in mixed methods. His latest book, "Revolution by Stealth: How Women's Groups Catalyzed a Cultural Transformation in Bihar" (with Shruti Majumdar and Paromita Sanyal), was published by Cambridge University Press in September 2026.

If you have comments on this article or wish to contribute to the discussion, please email letters@indevelopmentmag.com. Responses will be featured in a letters section.