

The Julius Silver, Roslyn S. Silver, and Enid Silver Winslow

DIALOGUES IN ARTS AND SCIENCE

ECONOMIC INEQUALITY, POLARIZATION AND CONFLICT

Debraj Ray

Silver Professor of Economics

September 2003



NEW YORK UNIVERSITY

Economic Inequality, Polarization and Conflict¹

In these notes, I attempt to outline and discuss questions of some importance in my field of research. In doing this, I am faced with a dilemma: I work, although not with frequent success, in more than one field. With some simplification, it is probably fair to say that I am currently pursuing two lines of research, one having to do with the study of income and wealth distribution, and the other addressing collective action and group formation. Providing an account of the many interesting questions in both these areas of research would make for a large, unwieldy and generally rambling document. So I have decided to restrict the scope of this essay to particular aspects of my research on income distribution. However, I invite you to visit my [homepage](#), which stores all my research, including working papers on the different topics that currently interest me.

To be sure, there is an enormous literature on issues of economic inequality. The purpose of these notes is to emphasize two questions that emerge from my own, necessarily narrower focus. In brief:

(a) does the market mechanism have some intrinsic, built-in drive towards the creation of economic inequality?

(b) what features of the distribution of income or wealth are likely to be closely correlated with the possibility of social unrest and conflict?

Phrased this way, my questions avoid significant philosophical issues on exactly what inequality means (though I will have something to say about this below). This literature concerns itself with the ethical meanings we might attach to an inequality measure, such as the Gini coefficient or the Theil index, and in this sense it represents an extension of the field of welfare economics. There is also a connection here to a literature even more replete with philosophical puzzles: one on what it even means to compare welfare across individuals, a requirement on which any meaningful definition of inequality must necessarily be based. [Amartya Sen's [*On Economic Inequality*](#) perhaps continues to be the best reference for understanding the ethical underpinnings of inequality measures.]

In contrast, my emphasis here is on the relationship between unequal outcomes and various aspects of societal structure or performance (social conflict being theme of interest in these notes). This is what one might call a functional viewpoint, complementary to but different from the “intrinsic question” of inequality per se. Thus in my research agenda, economic inequality is of interest, not only in its own right, but for its close ties to diverse economic and social variables.

1. The Evolution of Inequality

What determines the evolution and persistence of inequality in an economic system? One might distinguish two broad approaches to this problem, each consistent with a distinct view on the nature of economic inequality. The

first view — perhaps the dominant one in modern economics — is that inequality is the outcome of a constant battle between “convergence” and “luck.”²

“Luck” is the easier of these two concepts, so let’s get the notion of convergence out of the way. Of all the basic ideas that have held sway over mainstream economic theory, this is probably the most counterintuitive to the intelligent outsider. It states that poorer economic agents are likely to experience faster economic growth than their richer counterparts. This central underpinning of the theory of economic growth goes back at least to Robert Solow, and perhaps earlier. The basic idea is very simple indeed. Suppose that all production is carried out using capital and labor, and a constant fraction of produced income is saved. Then economic agents (countries, regions, even individuals) with a low endowment of capital relative to labor will have a high rate of return to capital (by the “law” of diminishing returns). Consequently, a given addition to the capital stock will have a larger marginal impact on that agent’s income. It follows that, controlling for savings rates, poorer agents will tend to grow faster and hence will catch up, converge.

Now the qualification “controlling for savings rates” is important. If the poor have lower savings rates than the rich, then the convergence process could be attenuated or even reversed. The same could happen if “agents” are replaced by “countries,” and poorer countries have higher population growth rates than their richer counterparts. Thus the convergence idea, properly interpreted, does not really mean that all economic do actually converge. But it does mean that a

failure to observe convergence must be traced to one or another of the additional factors that we've just described, such as varying savings rates. That could be one way to go (though so far as inequality is concerned it is not clear where this route would take us), but as I mentioned earlier, the traditional approach keeps in line with convergence, but considers frequent uncertain "shocks" to agent wealth. These two forces — convergence, and "shocks" or "luck" — pull in opposite directions, and the outcome of this tension determines the "equilibrium" distribution of income or wealth.

This approach to the evolution of inequality is naturally tied up with a particular view of inequality. In the ergodic distribution of this process, there is inequality, but there is also mobility (the degree of such mobility being given by both the frequency and ahistoricity of the shocks as well as the speed of convergence). This is not to say that the inequality is unimportant. After all, mobility will only yield dividends (or punishments) at some future, possibly distant date. But the mobility is there, nonetheless.

In contrast, a second approach to the evolution of inequality begins by questioning the basic assumptions that lie behind the convergence theorems of Solow and others. As we have seen, the diminishing-returns postulate is critical to the convergence idea. At the same time, many production activities that exhibit economies of scale, or large setup costs, do not suffer from diminishing returns; quite the opposite. The rate of return to investment may well increase with the level of investment. Ahistoricity, or the insensitivity to initial condi-

tions subsumed in the notion of convergence, breaks down as soon as these technologies enter the picture. There are many early formal statements of this observation. Perhaps the clearest analysis comes from Colin Clark's theoretical work on the management of fisheries.³ Regeneration functions for the growth of fish populations are typically "nonconvex" over some region, in that the net rate of stock growth will increase with the stock itself over some interval. So the path taken by the optimal harvesting program depends on starting conditions. If the initial stock of fish is low, extinction is the long-run outcome. Otherwise a positive steady state stock will be reached and maintained.

These ideas have natural applications to the theory of inequality. Imagine a population of individuals, each immersed in one of these nonconvex accumulation problems. Then historical disparities may persist and magnify, if initial differences among the population straddle some bifurcation threshold. Observe, however, that imperfections in credit markets are fundamental to the argument. With perfect credit markets, individuals would choose the best investment project regardless of initial conditions, and inequality would tend to disappear (essentially, the nonconvexity is removed by the introduction of a well-functioning credit market).

As with the first approach to inequality (based on the interplay between luck and convergence), this approach is also tied to a particular view of economic development. It is that current inequalities are reinforced by anterior inequalities (perhaps those prevailing in the generation of one's parents). The emphasis

is not on mobility or ergodicity. The emphasis is on persistence. This is not to say that these models are incompatible with substantial mobility. For example, the introduction of stochastic shocks into this scenario would certainly create a certain degree of flexibility. But the focus of this literature is different.

My joint research with Dilip Mookherjee (Boston University) falls into this second category, though the emphasis is not so much on economic agents accumulating wealth in isolation as on market interactions between them. One of our preliminary questions is: is inequality *inevitable*? Now at a certain level, it surely is: to the extent that there are random fluctuations in individual wealth or income, inequality is likely to arise. But the conceptual underpinning of our question is different: is there something basic about the way the market mechanism works that must cause inequality to emerge, with or without stochastic fluctuations? Phrased this way the answer seems to be an obvious no: if all individuals have identical preferences and start out identical, must they not make identical choices? For instance, even in the second approach described above, inequality persists once it arises. But the framework is also consistent with no inequality at all: if one starts from perfect equality, there will be no reason (absent stochastic shocks) for inequality to emerge. However, in reexamining this seemingly settled question, however, some interesting observations arise.

Consider the simplest example needed to express our idea: a world with just two occupations, or two kinds of human inputs in production, which we'll call skilled labor and unskilled labor. To acquire skills for each generation of

workers requires the upfront payment of an “educational cost,” while to remain unskilled cost requires no upfront payment. To complete the description of the example, suppose that firms combine skilled and unskilled labor to produce output. Now suppose that all individuals have identical wealth to begin with. Note, nevertheless, that *they cannot all take the same educational decision*. For if everyone in a new generation were to stay unskilled, the skilled wage would become exceedingly high, and the acquisition of skills would be highly rewarding. At the same time, a situation in which all individuals acquire skills is not sustainable either. For then the *unskilled* wage would be very high. [To be sure, if educated individuals can costlessly perform unskilled tasks, this would not happen, but in any case the unskilled wage would then become equal to the skilled wage, with exactly the same implications that we now discuss.] Neither of these two situations is, therefore, compatible with an “equilibrium outcome.”

In equilibrium, therefore, symmetry *must* be broken. Identical people don’t make identical choices: indeed, in this example, they cannot. To be sure, as far the current generation is concerned, this means nothing in terms of payoffs. If individuals have identical resources and preferences, they must receive the same payoffs even though they take different actions. However, while identical-wealth individuals are indifferent in equilibrium as to whether or not to educate their children, their *children* are not indifferent. They would prefer to have the skills rather than not have them. So, in the next generation, there *must* be individuals whose parents apparently scrimped and saved, so that their

offspring could lead a better life. And there must be individuals whose parents consumed all their income and did not educate their children. Of course, once such inequality is generated, it will magnify and persist. It is now easier for higher-income parents to acquire skills for the children, so the process will cumulate. This second part is more intuitive, what is perhaps less intuitive is that the market mechanism must cause such a process to inevitably start, even if we begin with all individuals perfectly equal.

Notice how the movement of prices (in our example, the changes in the skilled and unskilled wages in response to changes in the supply of the two different types of labor) are crucial in generating the initial broken symmetry and the consequent emergence of inequality. But this is only the starting point of our research, along with the work of several others. The framework itself is now ripe for the study of several related questions:

1. Might it be possible for an economy — the same economy — to display a multiplicity of stable steady states, each with a different degree of inequality? If so, under what conditions? One of our papers suggests that there is an intriguing connection between multiplicity and the richness of occupational structure. The richer that structure, the less is the tendency for multiplicity. Without dwelling on this connection in greater detail, let me highlight why the question of multiplicity is all-important. If such multiplicity exists, the very same economy could behave in dramatically different ways, depending on the starting point for that economy. Certain policy interventions would simply have a transient effect: it

would bring the economy back to the point where it started from. But others might engineer a move to an entirely different steady state, with significant consequences for the degree of inequality and for the values of associated economic variables. In the face of multiplicity, *judiciously chosen one-time interventions may have a lasting and permanent effect.*

2. Starting from arbitrary initial conditions, does the economy always converge to *some* steady state? This question is of importance because it will either justify or largely invalidate analysis based on steady states. For instance, suppose that a one-time intervention, such as a land reform, reduces wealth disparities. It seems only reasonable that if the economy does not return to its old steady state, it will transit to a new steady state with lower inequality. But this isn't infallible logic by any means. It requires us to study out-of-steady state behavior. It must be emphasized that this problem, apart from being very important, is also very hard. Future research will have to study carefully — and in more detail — the subtle and often complex connections between initial conditions and final steady states.

3. Are more unequal steady states associated with greater degrees of economic inefficiency? In the sort of model we study, the answer is “yes.” With imperfect capital markets, greater inequality translates into greater paucity of access to opportunities, generating greater inefficiency. But as one might imagine, the general question transcends the particular framework I have described. What if “higher-ability” people are best suited for the best jobs: isn't

the fact that there is inequality permit them to access those jobs with greater ease? Unfortunately, we have very little idea of the degree to which wealth and ability are correlated. Very roughly speaking, starting wealth (which will determine access) is correlated with parental ability, and we know little about inter-generational correlations in ability. This connection between inequality and efficiency deserves (and will receive) careful scrutiny for a variety of reasons.

2. Distribution and Conflict

It is, of course, a truism that differences (perceived or otherwise) are at the heart of conflictual situations, and that economic inequality — as a fundamental component of differences across individuals — will have much to do with conflict, social tensions, or rebellion. The economic costs of such conflict are fairly self-evident. In its most brutal form there is loss of life. But it would be easy enough to point out other costs in conflict situations that are not life-threatening: the waste of resources, human and material, in protests, rallies, curfews and strikes, as well as the possible misallocation of resources among different productive activities. It is no surprise, then, that the relationship between distribution and conflict is of special interest.

My joint research with Joan Esteban at the Instituto de Análisis Económico in Barcelona develops the concept of polarization as an important potential determinant of social conflict. Part of our research deals with the conceptualization and measurement of polarization, and argues that this variable should be consid-

ered separately from inequality. In fact, the purpose of this research project is to argue that “the phenomenon of polarization [as distinct from inequality] is closely linked to the generation of tensions, to the possibilities of articulated rebellion and revolt, and to the existence of social unrest in general.”²⁴

The previous paragraph, and especially the stress it lays on the difference between polarization and inequality, is understandably meaningless to the reader uninitiated in these terms. Let me, then, first clarify the definition of inequality that is standard in the economics literature. At the heart of the exercise is an attempt to rank various income or wealth distributions by the amount of inequality they display. The central axiom that underlies this is known as the transfers principle due to Pigou and Dalton, which states that if one distribution can be obtained from another by a sequence of regressive transfers from relatively poor to relatively rich, then the former should be deemed to be more unequal. (Observe: no actual sequence of transfers need connect the two distributions. It suffices that such transfers can in principle be shown to exist.) It should be fairly clear that every arbitrary pair of income distributions cannot be unambiguously compared in this fashion, in which the Pigou-Dalton criterion is silent. Mathematically, then, what the Pigou-Dalton axiom yields is a partial ordering on the space of income or wealth distributions. Every inequality measure used in economics (such as the coefficient of variation, the Gini coefficient, or the Theil entropy index) is compatible with this partial ordering, in that they rank any pair of distributions in the same way as Pigou-Dalton does, whenever the latter is applicable.

Now consider the idea of polarization introduced in my research with Esteban. At the risk of excessive narcissism, I quote heavily from our joint paper, because I don't think I can say it much better today:

“Consider a particular distribution [of characteristics, which could be income or any other measurable attribute]. Suppose that the population is grouped into significantly-sized ‘clusters,’” such that each cluster is very ‘similar’ in terms of the attributes of its members, but different clusters have members with very “dissimilar” attributes. In that case we would say that the society is ‘polarized’. Our purpose is to study these intuitive criteria carefully, and to provide a theory of measurement.

...We begin with the obvious question: why are we interested in polarization? It is our contention that the phenomenon of polarization is closely linked to the generation of tensions, to the possibilities of articulated rebellion and revolt, and to the existence of social unrest in general. This is especially true if the underlying set of attributes is a variable such as income or wealth. A society that is divided into groups, with substantial intra-group homogeneity and inter-group heterogeneity in, say, incomes, is likely to exhibit the features mentioned above. At the same time, measured inequality in such a society may be low.”⁵




Figure 1:
Polarization.

Let me try to illustrate the difference between polarization and inequality by means of an example. On the horizontal axis of Figure 1 above are various incomes ranging from low to high, and on the vertical axis we measure densities of the population earning different values of these incomes. Begin with the two relatively flat humps on the left and the right, and imagine that this represents the distribution of income in some hypothetical society. The distribution is then progressively “spiked” into two local modes, as indicated by the arrows: the population becomes located in more concentrated fashion on the two modal income values. (This, by the way, is not unlike the twin-peaks property identified for world income distributions by Danny Quah and others.⁶)

Which distributions exhibit greater “polarization”: the original, relatively flat distribution or its spikier cousin? We would argue for the latter. Two groups come into sharper focus, and each therefore has a stronger sense of itself and the other. Under the spikier distribution the population is more “rich-poor,” with no “middle class” bridging the gap between the two, and one may be inclined to perceive this situation as more conflictual than the initial one. If you are still hesitant, think of the attribute as an opinion index over a given political issue, running from right to left. Many would agree that political conflict is more likely under a two-spike distribution — with perhaps not completely extreme political opinions, but sharply defined and involving population groups of significant size — rather than under the flatter distribution.

But the point is this. If you admit the possibility of greater polarization in the spikier distribution, you are forced to depart from the domain of inequality

measurement. For under any inequality measure that is consistent with the widely accepted Pigou-Dalton axiom, inequality comes down as spikiness increases!

What creates this departure of polarization from inequality? What is fundamentally different is that polarization requires identification as well as alienation. There must be a sense of the “other,” something that sits well with inequality, but there must also be a sense of the “self,” a concept that is fed by increased equality (at least locally). Polarization, we argue, is a concept that has elements of both equality and inequality.

In my own development of the concept of polarization, first with Joan Esteban and more recently with Jean-Yves Duclos at the Université Laval in Quebec, I have relied almost exclusively on what might be called the identification-alienation frame- work (IA). To follow up on the idea in the last paragraph, polarization is related to the alienation that individuals and groups feel between one another, but such alienation is fueled by the degree of within-group identity. In concentrating on such phenomena, however, we do not mean to suggest that instances in which a single isolated individual runs amok with a machine gun are rare, or that they are unimportant in the larger scheme of things. Rather, these are not the objects of our enquiry. We are interested in the correlates of organized, large-scale social unrest.

The reader interested in the conceptual and technical details — and there are many — is warmly invited to study our [three-author paper](#). Here, I will restrict myself to an outline of how one approaches these issues. The basic

strategy is axiomatic. We begin by setting down a set of simple rules that any polarization measure must satisfy. Then we mathematically characterize what that class of measures must be. The advantage of this approach is that it lays bare the basic ingredients underlying the measure. The characterization implies that if we do not like how the measure behaves in a particular situation, we must, at some level, disagree with one or more of the axioms or the primitive structure of the model (in this case, the IA framework).

The domain under consideration is the class of all continuous distributions of income, with their integrals corresponding to various population sizes. Let F be such a distribution. An individual located at income x is presumed to feel a sense of identification that depends on the density of F at x , $f(x)$. At the same time, she feels alienation $|x-y|$ as far as an individual located at y is concerned. We write the effective antagonism of x towards y (under F) as some nonnegative nondecreasing function $T(i, a)$, where $i = f(x)$ and $a = |x - y|$. Finally, we take polarization to be proportional to the “sum” of all effective antagonisms:

$$P(F) = \iint T(f(x), |x - y|) f(x)f(y) dx dy.$$

But written in this way, the polarization measure is pretty useless, much depending on the choice of the functional form T . We therefore combine this approach with four axioms. The axioms are placed on densities that are unions of one or more unimodal, symmetric “basic densities.” Such densities will be scaled down and up to accommodate varying populations. They can be income-scaled, or moved from one side to another to capture uniform absolute income

changes. Finally a basic density can be squeezed, to move population mass away from the sides to the center. (See [Polarization: Concepts, Measures, Estimation](#) for a formal definition and a bit more discussion.)

Now consider the following four axioms (consult Fig. 2 for an illustration):

Axiom 1. “Single-Squeeze.” If a distribution is composed of a single basic density, then a squeeze of that basic density cannot increase polarization.




Figure 2: Illustration of the Axioms for Polarization

Axiom 2. “Double-Squeeze.” If a symmetric distribution is composed of three basic densities drawn from the same kernel, with mutually disjoint supports, then a symmetric squeeze of the side densities cannot reduce polarization. (Note: this is the central assumption that distinguishes polarization from inequality.)

Axiom 3. “Slide.” Consider a symmetric distribution composed of four basic densities drawn from the same kernel, with mutually disjoint supports. Slide the two middle densities to the side as shown (keeping all supports disjoint). Then polarization must go up.

Axiom 4. “Population-Homotheticity.” Identical population changes for two distributions cannot affect their polarization ranking.

Now one can prove the following characterization theorem:

A polarization measure satisfies the general class described above and Axioms 1–4 if and only if

$$P_{\alpha}(f) \equiv \int \int (f(x)^{1+\alpha} f(y) |y - x| dx dy,$$

where the number α must lie between 1/4 and 1.

The usefulness of this sort of characterization should be clear, even if you are not technically-minded and haven’t followed the mathematical details. It means that we can actually measure polarization for different countries or regions, just as we do for inequality. It opens the door to correlating these measures with indices of conflict or unrest, and moves us towards a potential understanding of what causes conflict. Research along these lines is already well under way. *Author’s update, March 2022: [Ethnicity and Conflict: Theory and Facts](#)* is an example.

Meanwhile, there is much work to do on the conceptual side as well. The measure described above is particularly adequate for polarizations of income distributions. However, we aren’t only interested in income. We may be interested in studying polarization across social, ethnic, religious, or geographically separated groups. Such groupings are not necessarily defined by income or wealth, though in a worldview where economic class is the only variable of

interest, it would be. (I hasten to add that I don't subscribe to that worldview.) So one item on the research agenda is to extend these measures to other dimensions of alienation or identification. Once again, at the risk of some unavoidable but fairly trivial mathematical notation, I indicate how this might be accomplished.

Suppose, then, that our basic groupings are not given by income, but by religious, caste, or other ethnic groupings, or groups based on criteria such as gender or spatial location. (Under what conditions certain grouping acquire social salience is also a question of immense interest, but I won't comment on that here.) So society is divided into G groups, the choice of which is left to the researcher. Continue to maintain the IA framework: each individual feels identification with those he considers to be members of his "own group," and alienation from those he considers to be members of "other groups." The form of the measure suggested by our characterization theorem can then be readily transplanted to "social polarization." Consider, for instance, the case of "pure social polarization," in which income plays no role. That is, assume each person is fully identified with every other member of his group. Likewise, the alienation function takes on values that are specific to group pairs $\{j, k\}$ and have no reference to income. Then, mimicking our derived equation,

$$P_s(\mathbf{F}) = \sum_{j=1}^G \sum_{k=1}^G n_j^\alpha n_k \Delta_{jk}$$

would represent social polarization in this case (where the Deltas represent "social" distances across groups, and the n_j 's are population weights within groups).

Once the two extremes — pure income polarization and pure social polarization — are identified, we may easily consider several hybrids. As examples, consider the case in which notions of identification are mediated not just by group membership but by income similarities as well, while the antagonism equation remains untouched. Then we get what one might call social polarization with income-mediated identification:

$$P_s(\mathbf{F}) = \sum_{j=1}^G (1 - n_j) \int_x f_j(x)^\alpha dF_j(x).$$

where the n_j 's are group population sizes as before and the f_j 's are the (unnormalized) income densities for each group. Indeed, one could expand (or contract) the importance of income further, while still staying away from the extremes. For instance, suppose that — in addition to the income-mediation of group identity — alienation is also income-mediated (that is, two individuals alienated from each other must belong to different groups and have different incomes). Now groups have only a demarcating role — they are necessary (but not sufficient) for identity, and they are necessary (but not sufficient) for alienation. The resulting measure would look like this:

$$P^*(\mathbf{F}) = \sum_{j=1}^G \sum_{k \neq j}^G \int_x \int_y f_j(x)^\alpha |x - y| dF_j(x) dF_k(y),$$

and so on. Note that we do not intend to suggest that other special cases or hybrids are not possible, or that they are less important. The discussion here is only to indicate that social and economic considerations can be profitably

combined in the measurement of polarization. One specific goal of our research project is to provide an axiomatic treatment of these hybrid measures.

3. Conflict in India

I have been engaged in theoretical research on the subject of conflict and its relationship with distribution; see, for instance, [Conflict and Distribution](#). However, I'd like to end these thoughts on a more applied note, with a discussion of the Indian case. What follows is preliminary to a fault, and is more a description of where I'd like to go rather than a well-thought-out research description. But it may be interesting enough to merit inclusion in these dialogues. The observations I will make are based entirely on forthcoming joint research with Anirban Mitra, now at the University of Kent. *Author's Update, March 2022*. This research is now published as [Implication of an Economic Theory of Conflict](#). However, I have left the paragraphs that follow entirely untouched, and they are presented here as written in 2003.

We plan to examine how changes in economic inequality, both within and across religious and ethnic groups, affects the likelihood of inter-group conflict. Our empirical research will be exclusively based on the Indian experience.

A brief background: Hindus, at 82% of the Indian population, constitute the religious majority, and Muslims, at 12%, the largest religious minority group. The Hindu nationalist narrative (*Hindutva*) argues that Hinduism is not just the religion of a majority of India's population, but that it defines India's national

identity. While other religions may exist in India, the *Hindutva* narrative clearly assigns them a subordinate role. Indeed, there has been a clear ascendancy of this idea in Indian politics. The principal patron of Hindu nationalism in mainstream politics is the Bharatiya Janata Party (BJP). The 1989 national elections marked the entry of the BJP as a major political player, and today it is the single largest party in the national ruling coalition. In contrast, no Muslim party is an important player in national politics.

The Hindu-Muslim split is not the only major social divide in India. Within the Hindus the caste system has led to an inequitable distribution of resources. Individuals belonging to lower castes were traditionally relegated to menial occupations and prevented from owning assets. These castes, commonly known as “scheduled castes,” constitute 16% of the Indian population. While economic progress amongst the members of this group remains limited, low caste political activism has witnessed a dramatic rise in the last two decades. A political party centered on, and led by, low caste individuals has been the main ruling party in the most populous Indian state (Uttar Pradesh) for much of the 1990s, and is an important member of the national ruling coalition. The last decade and a half has also seen large increases in the size and scope of affirmative action programs directed towards these groups. (At Independence, the Indian government made a commitment to pursue affirmative action policies in favor of historically disadvantaged castes.)

These two inequities — Hindu-Muslim and high caste/low caste — are savagely mirrored in the escalation of group violence in India since the 1980s.

For instance, the average number of Muslims killed or injured in religious violence in India over this period is four times the average pre-1980 figures. Much of this violence has been related to a quasi-symbolic goal of the Hindu nationalist movement: to demolish a mosque (the Babri Masjid in Ayodhya), and replace it by a temple.

Moreover, this upsurge has a distinct regional character. According to the estimates of Steven Wilkinson, the average number of riots per month between 1970 and 1995 varied from a high of 0.80 per month in the Western Indian state of Gujarat (followed by Maharashtra at 0.54) to 0.03 per month in the Southern state of Kerala. The most recent upsurge of violence in Gujarat (February 2002) saw over 2,000 Muslims killed and 100,000 displaced. The organized nature of this violence, and the urban nature of such violence (with the targeting of Muslim businesses) has received an enormous amount of attention, even in the American press.

Caste violence in India also tends to have a regional flavor. Such conflict has been closely related to land disputes, and concentrated in states with a history of leftist guerrilla organizations. These states include West Bengal, Bihar and Andhra Pradesh. While low-caste individuals have long been victims of upper-caste violence, a recent trend has been towards organized conflict between high- and low- caste private armies.

Finally, the last two decades have been a period of sustained economic growth in India. After averaging 3.6% annual GDP growth rate between 1950–51

and 1980–81, GDP growth accelerated to 5.6% in the 1980s and an even higher 6% in the 1990s. [According to the National Account Statistics, this growth is mirrored in a corresponding growth in consumption per capita, and a corresponding decline in overall poverty.] There is some evidence that this income growth (especially in the last decade) has been accompanied by an increase in income inequality. However, there are no estimates of how these income gains of the last two decades have been distributed across major religious and ethnic groups.

How does one bring questions of economic well-being into issues of religious or caste conflict? To be sure, we are not the first to ask this sort of question, though we believe that our starting point is novel (yet very simple). This has to do with the identity of the group that initiates conflict. To illustrate, suppose that there are two groups, *M* and *H*. In all cases we will think of group *M* as being poorer, on average, than group *H*. But we will entertain different hypotheses regarding the group that can politically initiate conflict.

View conflict as a battle for resources in a resource-scarce economy. For instance, one group may agitate for land reform if it is felt that the other group owns a lot of land. Or the business community in one group may feel threatened if there is substantial expansion in business interests among the rival group. Begin with the assumption that it is politically infeasible for *M* to initiate conflict, while it is politically feasible for *H* to do so [We will discuss reasons for this below]. The fundamental asymmetry is that *H* is richer, so that under this political feasibility assumption, we have

Conjecture 1. If H initiates conflict, conflict will be greater the greater the economic equality between the two groups.

What drives this conjecture is a simple notion of perceived threat. A richer group will feel more threatened if a poorer group comes closer to it in economic space, if that group encroaches on privileges and opportunities that the rich formerly had. In such a world, greater inequality between the two groups is actually more conducive to inter-group calm.

Related to the assumption that the initiator of conflict is the richer group is a conjectured description of just how such conflict is carried out. As exemplified by recent events in Gujarat, the richer group will typically compensate (or subsidize) other individuals to carry out the physical acts of demolition, intimidation, looting and violence against their poorer counterpart. If such “conflict labor” can only be drawn from one’s own group, we may derive

Conjecture 2. Holding all else constant, greater within- H inequality will raise the likelihood of conflict, because it makes it easier to buy and sell “conflict labor.”

These conjectures stand in explicit contrast to the models of socio-economic polarization discussed earlier, which emphasize the role of low within-group inequality (greater group identification) and high across-group inequality (greater cross-group alienation) in creating social conflict. Conjectures 1 and 2 reflect, simply, the opposite of these postulates. Matters are reversed, however, when we flip the conflict initiator.

Now suppose that it is politically feasible for M , but not for H , to initiate conflict. Then M feels more deprived as the gap between H and M widens, so that Conjecture 1 is amended to:

Conjecture 1A. If M initiates conflict, conflict will be greater the greater the economic inequality between the two groups.

Indeed, there is a sense in which Conjecture 2 tends to be overturned as well. A poor group, simply by virtue of its poverty, is less likely to employ conflict labor. Each individual must participate in the demonstration or in the act of rebellion. This makes conflict more likely when the poor group is more cohesive, and its members feels greater identification with one another:

Conjecture 2A. Holding all else constant, greater within- M inequality will raise the likelihood of conflict.

Conjectures 1A and 2A, taken together, have very different implications from their counterparts 1 and 2, much more in line, for instance, with the standard view of polarization. Yet these sharply contrasting effects (which we plan to model formally in our research) arise from one very simple yet critical question: who initiates conflict? To understand this question in the context of India, which is the focus of this project, one must return to questions of group identity. In India, it would be unambiguously disadvantageous for Muslim groups to engage in conflict against Hindus. There are three reasons for this: first, the Muslims in India constitute a minority and such conflict would, in all likeli-

hood, lead to extraordinary repercussions being visited on the Muslims. Second, Muslim-initiated conflict would be viewed as profoundly anti-national, especially now that the notion of national identity has increasingly become equated with Hindu identity. Finally, to the extent that the bureaucracy, the police and the politicians are increasingly “Hindu-ized,” one would expect Muslim initiation of conflict to be immediately put down.

In contrast — and partly for the same reasons — it has become increasingly easy for Hindus to engage in acts of violence against the Muslims. This is why we maintain the assumption that *H* initiates conflict when *H* stands for Hindu and *M* stands for Muslim, and our main claim is that Conjectures 1 and 2 will hold in this context. To be sure, these conjectures need to be empirically studied, which is the object of our research.

Now switch to an entirely different context, in which *M* stands for the “menials” — the low-caste groups — and *H* stands for the high caste. Now it is possible to argue that the political and cultural rules are reversed. The Indian independence movement has long been associated with rights for low castes. A continuing cry for such rights — such as land reform — lies well within the Indian nationalist narrative, and can be tolerated for this reason. It may also be politically easier: half a century of affirmative action has ensured that a growing number of low-caste individuals occupy political or bureaucratic power, and political power among the low castes is by no means insignificant. Thus in this scenario, where *H* stands for “high caste” and *M* stands for the “menials,” our claim is that Conjectures 1A and 2A will apply.

The analysis so far pushes government policy well into the background, which is undesirable. In the list of “other things held equal” is the nature of State or district politics, which in turn create an atmosphere of tolerance for, or resistance to, conflict of various forms. One may instead posit that government policy is responsive to conflict. To examine this possibility, we will consider a government who is motivated by electoral concerns, and can choose to respond to conflict situations with a policy favorable to one of the two religious groups. Such policy may take the form of clamping down on any form of intergroup conflict, or instead turning a blind eye to ongoing violence. If the government has complete information about the demographic composition of the polity and their political preferences we would expect the government response to favor the median voter. Hence, policy will vary with the demographic composition of the polity in a well-defined manner. If, instead, the government has incomplete information then a group’s decision to engage in conflict potentially carries information about the distribution of preferences in the population. This, in turn, may alter the observed relationship between demographics and public policy.

Our research project will attempt to carefully conceptualize and draw out the implications of these conjectures.

4. Conclusions

I’ve tried here to document some of my own interests in the area of inequality, polarization, and conflict. These notes are very self-centered — they do not enumerate all the research questions in this area (of which there are

surely many), but only my relatively narrow agenda. Yet, I hope that the sample of issues taken up here will serve as an illuminating entry point into this fascinating literature. I also hope that when I am next given the opportunity to revise these notes, I will have new results to report and new directions to point to. For now, here is a summary. I have noted three areas of research, some more theoretical than others:

1. Is there a tendency for market economies to generate inequality? This is an important question, not because inequality may be viewed as problematic in its own right (which you and I may or may not agree on), but because persistent inequality has its own effects on social outcomes. One such outcome is social unrest or conflict, which motivates my second line of inquiry:

2. What, precisely, are those characteristics of economic distribution that are most strongly correlated with potential conflict? In this section, I've argued that one must investigate the notion of polarization, which I attempt to describe.

3. Finally, and as an empirical counterpart to the relatively abstract ideas in (2), I discuss violence in India and the possible economic correlates of such violence. As I've indicated, however, this is very preliminary research, and most of the (exciting) work lies ahead.

1 *Author's Note, March 2022.* These notes were written in 2003 as part of the “Silver Dialogues” by incoming Silver Professors at NYU. I was fortunate to have received this honor in 2002. I haven’t updated the notes at all, but do want to add that the projects described here have since been built upon to greater or lesser degree. My [homepage](#) contains these publications, and you are warmly invited to visit and explore, or to get in touch by emailing me: debraj.ray@nyu.edu.

2 See, for instance, Gary Becker and Nigel Tomes (1979), “An Equilibrium Theory of the Distribution of Income and Intergenerational Mobility,” *Journal of Political Economy* 87, 1153–1189.

3 Colin Clark (1971). “Economically Optimal Policies for the Utilization of Biologically Renewable Resources,” *Mathematical Biosciences* 12, 245–260.

4 Joan Esteban and Debraj Ray (1994), “On the Measurement of Polarization,” *Econometrica* 62, 819–851.

5 Esteban and Ray, *Ibid.*

6 See, for instance, Danny Quah (1993), “Empirical Cross-Section Dynamics in Economic Growth,” *European Economic Review* 37, 426–434.