

# Catastrophe, Social Collapse, and Human Extinction

Robin Hanson\*  
Department of Economics  
George Mason University†

August 2007

## Abstract

Humans have slowly built more productive societies by slowly acquiring various kinds of capital, and by carefully matching them to each other. Because disruptions can disturb this careful matching, and discourage social coordination, large disruptions can cause a “social collapse,” i.e., a reduction in productivity out of proportion to the disruption. For many types of disasters, severity seems to follow a power law distribution. For some of types, such as wars and earthquakes, most of the expected harm is predicted to occur in extreme events, which kill most people on Earth. So if we are willing to worry about any war or earthquake, we should worry especially about extreme versions. If individuals varied little in their resistance to such disruptions, events a little stronger than extreme ones would eliminate humanity, and our only hope would be to prevent such events. If individuals vary a lot in their resistance, however, then it may pay to increase the variance in such resistance, such as by creating special sanctuaries from which the few remaining humans could rebuild society.

## Introduction

“Modern society is a bicycle, with economic growth being the forward momentum that keeps the wheels spinning. As long as the wheels of a bicycle are spinning rapidly, it is a very stable vehicle indeed. But, [Friedman] argues, when the wheels stop - even as the result of economic stagnation, rather than a downturn or a depression - political democracy, individual liberty, and social tolerance are then greatly at risk even in countries where the absolute level of material prosperity remains high....” (DeLong, 2006)

---

\*For their comments I thank Jason Matheny, the editors, and an anonymous referee. For their financial support, I thank the Center for Study of Public Choice and the Mercatus Center.

†[rhanson@gmu.edu](mailto:rhanson@gmu.edu) <http://hanson.gmu.edu> 703-993-2326 FAX: 703-993-2323 MS 1D3, Carow Hall, Fairfax VA 22030

The main reason to be careful when you walk up a flight of stairs is not that you might slip and have to retrace one step, but rather that the first slip might cause a second slip, and so on until you fall dozens of steps and break your neck. Similarly we are concerned about the sorts of catastrophes explored in this book not only because of their terrible direct effects, but also because they may induce an even more damaging collapse of our economic and social systems. In this chapter, I consider the nature of societies, the nature of social collapse, the distribution of disasters which might induce social collapse, and possible strategies for limiting the extent and harm of such collapse.

## What Is Society

Before we can understand how societies collapse, we must first understand how societies exist and grow. Humans are far more numerous, capable, and rich than were our distant ancestors. How is this possible? One answer is that today we have more of most kinds of “capital,” but by itself this answer tells us little; after all, “capital” is just anything that helps us to produce or achieve more. We can understand better by considering the various types of capital we have.

First we have natural capital, such as soil to farm, ores to mine, trees to cut, water to drink, animals to domesticate, and so on. Second, we have physical capital, such as cleared land to farm, irrigation ditches to move water, buildings to live in, tools to use, machines to run, and so on. Third we have human capital, such as healthy hands to work with, skills we have honed with practice, useful techniques we have discovered, and abstract principles that help us think.

Fourth we have social capital, i.e., ways in which groups of people have found to coordinate their activities. For example, households organize who does what chores, firms organize which employees do which tasks, networks of firms organize to supply inputs to each other, cities and nations organize to put different activities in different locations, culture organizes our expectations about the ways we treat each other, law organizes our coalitions to settle small disputes, and governments coordinate our largest disputes.

There are several important things to understand about all this capital. First, the value of almost any piece of capital depends greatly on what other kinds of capital are available nearby. A fence may be very useful in a prairie but useless in a jungle, while a nuclear engineer’s skills may be worth millions in a rich nation, but nothing in a poor nation. The productivity of an unskilled laborer depends greatly on how many other such laborers are available.

Second, scale makes a huge difference. The more people in a city or nation, the more each person or group can narrow their specialty, and get better at it. Special products or services that would just not be possible in a small society can thrive in a large society. So anything that lets people live more densely, or lets them talk or travel more easily, can create large gains by increasing the effective social scale.

Third, coordination and balance of capital are very important. For example, places with low social capital can stay poor even after outsiders contribute huge resources and training,

while places with high social capital can quickly recover from wars that devastate their natural, physical, and human capital.

## Social Growth

The opposite of collapse is growth. Over history we have dramatically increased our quantities of most, though not all, kinds of capital. How has this been possible?

Over the last few decades economists have learned a lot about how societies grow (Barro & Sala-I-Martin, 2003; Jones, 2002; Aghion & Howitt, 1998). While much ignorance remains, a few things seem clear. Social capital is crucial; rich places can grow fast while poor places decline. Also crucial is scale and neighboring social activity; we each benefit greatly on average from other productive activity nearby.

Another key point is that better “technology,” i.e., better techniques and coordination, drive growth more than increased natural or physical capital. Better technology helps us produce and maintain more natural and physical capital, a stronger effect than the ability of more natural and physical capital to enable better technology (Grubler, 1998).

To complete our mental picture, let us quickly review the history of growth (Hanson, 2000), starting with animals. All animal species have capital in the form of a set of healthy individuals and a carefully honed genetic design. An individual animal may also have capital in the form of a lair, a defended territory, and experience with that area. Social animals, such as ants, also have capital in the form of stable organized groups.

Over many millions of years the genetic designs of animals slowly acquired more possibilities. For example, over the last half billion years the size of the largest brains doubled roughly every thirty five million years. About two million years ago some primates acquired the combination of a large social brain, hands that could handle tools, and mouths that could voice words; a combination that allowed tools, techniques, and culture to become powerful forms of capital.

The initial human species had perhaps ten thousand members, which some estimate to be the minimum for a functioning sexual species. As human hunter-gatherers slowly accumulated more kinds of tools, clothes, and skills, they were able to live in more kinds of places, and their number doubled every quarter million years. Eventually, about ten thousand years ago, humans in some places knew enough about how to encourage local plants and animals that these humans could stop wandering and stay in one place.

Non-wandering farmers could invest more profitably in physical capital such as cleared land, irrigation ditches, buildings, and so on. The increase in density that farming allowed also enabled our ancestors to interact and coordinate with more people. While a hunter-gatherer might not meet more than a few hundred people in his or her life, a farmer could meet and trade with many thousands.

Soon, however, these farming advantages of scale and physical capital reached diminishing returns, as the total productivity of a region was limited by its land area and the kinds of plants and animals available to grow. Growth was then limited importantly by the rate at which humans could domesticate new kinds of plants and animals, allowing the colonization

of new land. Since farmers talked more, they could spread such innovations much faster than hunter-gatherers; the farming population doubled every thousand years.

A few centuries ago, the steady increase in farming efficiency and density, as well as travel ease, finally allowed humans to specialize enough to support an industrial society. Specialized machines, factories, and new forms of social coordinate, allowed a huge increase in productivity. Diminishing returns quickly set in regarding the mass of machines we produced, however. We still make about the same mass of items per person as we did two centuries ago.

Today's machines are far more capable though, because of improving technologies. And networks of communication between specialists in particular techniques have allowed the rapid exchange of innovations; during the industrial era world product (the value of items and services we produce) has doubled roughly every fifteen years.

Our history has thus seen four key growth modes; animals with larger brains, human hunter-gatherers with more tools and culture enabling them to fill more niches, human farmers domesticating more plants, animals, and land types, and human industry improving its techniques and social capital. During each mode growth was over a hundred times faster than before, and production grew by a factor of over two hundred. While it is interesting to consider whether even faster growth modes might appear in the future, in this chapter we turn our attention to the opposite of growth: collapse.

## Social Collapse

Social productivity fluctuates constantly in response to various disturbances, such as changes in weather, technology or politics. Most such disturbances are small, and so induce only minor social changes, but the few largest disturbances can induce great social change. The historical record shows at least a few occasions where social productivity fell rapidly by a large enough degree to be worthy of the phrase "social collapse." For example, there have been famous and dramatic declines, with varying speeds, among ancient Sumeria, the Roman empire, and the Pueblo peoples. A century of reduced rain, including three droughts, apparently drove the Mayans from their cities and dramatically reduced their population, even though the Maya had great expertise and experience with irrigation and droughts (?; Haug, Gnther, Peterson, Sigman, Hughen, & Aeschlimann, 2003).

Some have explained these historical episodes of collapse as due to a predictable internal tendency of societies to overshoot ecological capacity (Diamond, 2005), or to create top-heavy social structures (Tainter, 1988). Other analysis, however, suggests that most known ancient collapses were initiated by external climate change (Weiss & Bradley, 2001; deMenocal, 2001). The magnitude of the social impact, however, often seems out of proportion to the external disturbance. Similarly in recent years relatively minor external problems often translate into much larger reductions in economic growth (Rodrik, 1999). This disproportionate response is of great concern; what causes it?

One obvious explanation is that the intricate coordination that makes a society more productive also makes it more vulnerable to disruptions. For example, productivity in our

society requires continued inputs from a large number of specialized systems, such as for electricity, water, food, heat, transportation, communication, medicine, defense, training, and sewage. Failure of any one of these systems for an extended period can destroy the entire system. And since geographic regions often specialize in supplying particular inputs, disruption of one geographic region can have a disproportionate effect on a larger society. Transportation disruptions can also reduce the benefits of scale societies enjoy.

Capital that is normally carefully balanced can become unbalanced during a crisis. For example, a hurricane may suddenly increase the value of gas, wood, and fresh water relative to other goods. The sudden change in the relative value of different kinds of capital produces inequality, i.e., big winners and losers, and envy, i.e., a feeling that winner gains are undeserved. Such envy can encourage theft and prevent ordinary social institutions from functioning; consider the widespread resistance to letting market prices rise to allocate gas or water during a crisis.

“End game” issues can also dilute reputational incentives in severe situations. A great deal of social coordination and cooperation is possible today because the future looms large. We forgo direct personal benefits now for fear that others might learn later of such actions and avoid us as associates. For most of us, the short term benefits of “defection” seem small compared to the long term benefits of continued social “cooperation.”

But in the context of a severe crisis, the current benefits of defection can loom larger. So not only should there be more personal grabs, but the expectation of such grabs should reduce social coordination. For example, a judge who would not normally consider taking a bribe may do so when his life is at stake, allowing others to expect to get away with theft more easily, which leads still others to avoid making investments that might be stolen, and so on. Also, people may be reluctant to trust bank accounts or even paper money, preventing those institutions from functioning.

Such multiplier effects of social collapse can induce social elites to try to deceive the rest about the magnitude of any given disruption. But the rest of society will anticipate such deception, making it hard for social elites to accurately communicate the magnitude of any given disruption. This will force individuals to attend more to their private clues, and lead to less social coordination in dealing with disruptions.

The detailed paths of social collapse depend a great deal on the type of initial disruption and the kind of society disrupted. Rather than explore these many details, let us see how far we can get thinking in general about social collapse due to large social disruptions.

## The Distribution of Disaster

First, let us consider some general features of the kinds of events that can trigger large social disruptions. We have in mind events like earthquakes, hurricanes, plagues, wars, revolutions, and so on. Each such catastrophic event can be described by its *severity*, which might be defined in terms of energy released, deaths induced, and so on.

For many kinds of catastrophes, the distribution of event severity appears to follow a power law over a wide severity range. That is, sometimes the chance that within a small

time interval one will see an event with severity  $S$  that is greater than a threshold  $s$  is given by

$$P(S > s) = ks^{-\alpha}, \quad (1)$$

where  $k$  is a constant and  $\alpha$  is the *power* of this type of disaster.

Now we should keep in mind that these powers  $\alpha$  can only be known to apply within the scales sampled by available data, and that many have disputed how widely such power laws apply (Bilham, 2004), and whether power laws are the best model form, compared for example to the lognormal distribution (Clauset, Shalizi, & Newman, 2007a).

Addressing such disputes is beyond the scope of this chapter. We will instead consider power law distributed disasters as an analysis reference case. Our conclusions would apply directly to types of disasters that continue to be distributed as a power law even up to very large severity. Compare to this reference case, we should worry less about types of disasters whose frequency of very large events is below a power law, and more about types of disasters whose frequency is greater.

The higher the power  $\alpha$ , the fewer larger disasters there are relative to small disasters. For example, if they followed a power law then car accidents would have a high power, as most accidents involve only one or two cars, and very few accidents involve one hundred or more cars. Supernovae deaths, on the other hand, would probably have a small power; if anyone on Earth is killed by a supernova, most likely many will be killed.

Disasters with a power of one are right in the middle, with both small and large disasters being important. For example, the energy of earthquakes, asteroid impacts, and Pacific hurricanes all seem to be distributed with a power of about one (Christensen, Danon, Scanlon, & Bak, 2002; Lay & Wallace, 1995; Morrison, Harris, Sommer, Chapman, & Carusi, 2003; Sanders, 2005). (The land area disrupted by an earthquake also seems to have a power of one (Turcotte, 1999).) This implies that for any given earthquake energy  $E$  and for any time interval, as much energy will on average be released in earthquakes with energies in the range from  $E$  to  $2E$  as in earthquakes with energies in the range from  $E/2$  to  $E$ . While there should be twice as many events in the second range, each event should only release half as much energy.

Disasters with a high power are not very relevant for social collapse, as they have little chance of being large. So, assuming published power estimates are reliable and that the future repeats the past, we can set aside windstorms (energy power of 12), and worry only somewhat about floods tornadoes, and terrorist attacks (with death powers of 1.35, 1.4, and 1.4). But we should worry more about disasters with lower powers, such as forest fires (area power of 0.66), hurricanes (dollar loss power of 0.98, death power of 0.58), earthquakes (energy power of 1, dollar loss and death powers of 0.41), wars (death power of 0.41), and plagues (death power of 0.26 for Whooping Cough and Measles) (Barton & Nishenko, 1997; Cederman, 2003; Turcotte, 1999; Sanders, 2005; Watts, Muhamad, Medina, & Dodds, 2005; Clauset, Young, & Gleditsch, 2007b; Rhodes, Jensen, & Anderson, 1997; Nishenko & Barton, 1995).

Note that energy power tends to be higher than economic loss power, which tends to

be higher than death power. This says that compared to the social loss produced by a small disturbance, the loss produced by a large disturbance seems out of proportion to the disturbance, an effect that is especially strong for disasters that threaten lives and not just property. This may (but does not necessarily) reflect the disproportionate social collapse that large disasters induce.

For a type of disaster where damage is distributed with a power below one, if we are willing to spend time and effort to prevent and respond to small events, which hurt only a few people, we should be willing to spend far more to prevent and respond to very large events, which would hurt a large fraction of the Earth's population. This is because while large events are less likely, their enormous damage more than makes up for their low frequency. If our power law description is not misleading for very large events, then in terms of expected deaths, most of the deaths from war, earthquakes, hurricanes, and plagues occur in the very largest such events, which kill a large fraction of the world's population. And those deaths seem to be disproportionately due to social collapse, rather than the direct effect of the disturbance.

## Existential Disasters

How much should we worry about even larger disasters, triggered by disruptions several times stronger than ones that can kill a large fraction of humanity? Well, if we only cared about the expected number of people killed due to an event, then we would not care that much whether 99% or 99.9% of the population was killed. In this case, for low power disasters we would care the most about events large enough to kill roughly half of the population; our concern would fall away slowly as we considered smaller events, and fall away quickly as we considered larger events.

A disaster large enough to kill off humanity, however, should be of special concern. Such a disaster would prevent the existence of all future generations of humanity. Of course it is possible that humanity was about to end in any case, and it is also possible that without humans within a few million years some other mammal species on Earth would evolve to produce a society we would respect. Nevertheless, since it is also possible that neither of these things would happen, the complete destruction of humanity must be considered a great harm, above and beyond the number of humans killed in such an event.

It seems that groups of about seventy people colonized both Polynesia and the New World (Murray-McIntosh, Scrimshaw, Hatfield, & Penny, 1998; Hey, 2005). So let us assume, as a reference point for analysis, that the survival of humanity requires that one hundred humans remain, relatively close to one another, after a disruption and its resulting social collapse. With a healthy enough environment, one hundred connected humans might successfully adopt a hunter-gatherer lifestyle. If they were in close enough contact, and had enough resources to help them through a transition period, they might maintain a sufficiently diverse gene pool, and slowly increase their capabilities until they could support farming.

Once they could communicate to share innovations and grow at the rate that our farming ancestors grew, humanity should return to our population and productivity level within twenty thousand years. (The fact that we have used up some natural resources this time

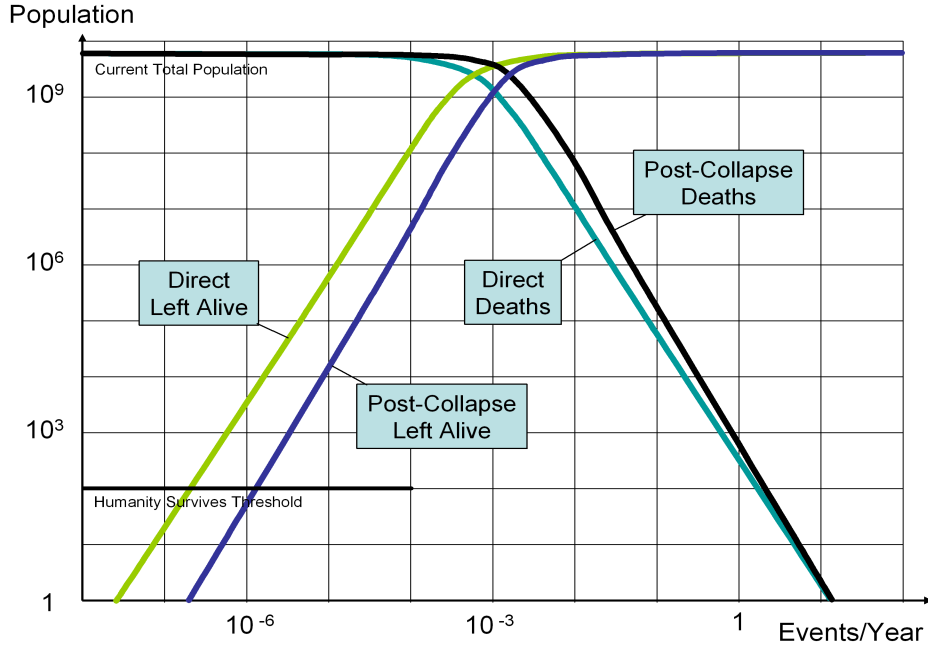


Figure 1: A Soft Cutoff Power Law Scenario

around would probably matter little, as growth rates do not seem to depend much on natural resource availability. ) With less than one hundred survivors near each other, on the other hand, we assume humanity would become extinct within a few generations.

Figure 1 illustrates a concrete example to help us explore some issues regarding existential disruptions and social collapse. It shows a log-log graph of event severity versus event frequency. For the line marked “Post-Fall Deaths,” the part of the line on the right side of the figure is set to be roughly the power law observed for war deaths today (earthquake deaths have the same slope, but are 1/3 as frequent). The line below it, marked “Direct Deaths” is speculative, and represents the idea that a disruption only directly causes some deaths; the rest are due to social collapse following a disruption. The additional deaths due to social collapse are a small correction for small events, and become a larger correction for larger events.

Of course the data to which these power laws have been fit does not include events where most of humanity was destroyed. So in the absence of direct data, we must make guesses about how to project the power law into the regime where most people are killed. If  $S$  is the severity of a disaster, to which a power law applies,  $T$  is the total population just before the disaster, and  $D$  is the number killed by the disaster, then one simple approach would be to set

$$D = \max(T, S). \quad (2)$$

This would produce a very hard cutoff.



In this case, either much of the population would be left alive or everyone would be dead; there would be little chance of anything close to the borderline. This model expresses the idea that whether a person dies from a disaster depends primarily on the strength of that disaster, and depends little on a varying individual ability to resist disaster. Given the parameters of figure 1, there would be a roughly a one in a thousand chance each year of seeing an event that destroyed all of humanity.

Figure 1 instead shows a smother projection, with a softer cutoff,

$$\frac{1}{D} = \frac{1}{S} + \frac{1}{T}. \quad (3)$$

In the regime where most people are left alive,  $D \ll T$ , this gives  $D \approx S$ , and so gives the familiar power law,

$$P(D > s) = ks^{-\alpha}. \quad (4)$$

But in the regime where the number of people left alive,  $L = T - D$ , is small, with  $L \ll T$ , we have a new but similar power law,

$$P(L < s) = k's^\alpha. \quad (5)$$

For this projection, it takes a much stronger event to destroy all of humanity.

This model expresses the idea that in addition to the strength of the disaster, variations in individual ability to resist disaster are also very important. Such power law survival fractions have been seen in some biological cases (Burchell, Mann, & Bunch, 2004). Variable resistance might be due to variations in geographic distance, stockpiled wealth, intelligence, health, and military strength.

Figure 1 shows a less than one in three million chance per year of an event that would kill everyone in the ensuing social collapse. But there is a one in five hundred thousand chance of an event that leaves less than one hundred people alive; by assumption, this would not be enough to save humanity. And if the remaining survivors were not all in one place, but distributed widely across the Earth and unable to move to come together, it might take many thousands of survivors to save humanity.

Figure 1 illustrates some of the kinds of tradeoffs involved in preventing the extinction of humanity. We assumed somewhat arbitrarily above that one hundred humans were required to preserve humanity. Whatever this number is, if it could be reduced somehow by a factor of two, for the survival of humanity that would be equivalent to making this type of disaster a factor of two less damaging, or increasing our current human population by a factor of two. In figure 1 that is equivalent to about a 25% reduction in rate at which this type of event occurs. This figure also predicts that of every fifty people left alive directly after the disruption, only one remains alive after the ensuing social collapse. A factor of two improvement in the number who survive social collapse would also bring the same benefits.

## Disaster Policy

For some types of disasters, like car accidents and windstorms, frequency falls so quickly with event severity that large events can be ignored; they just don't happen. For other types of disasters, such as floods, tornadoes, and terrorist attacks, the frequency falls quickly enough that disasters large enough to cause serious social collapse can be mostly ignored; they are very rare.

But for still other types of disasters, such as fires, hurricanes, earthquakes, wars, and plagues, most of the expected harm may be in the infrequent but largest events, which would hurt a large fraction of the world. So if we are willing to invest at all in preventing or preparing for these type of events, it seems we should invest the most in preventing and prepare for these largest events. (Of course this conclusion is muted if there are other benefits of preparing for smaller events, benefits which do not similarly apply to preparing for large events.)

For some types of events, such as wars or plagues, large events often arise from small events that go wrong, and so preparing for and preventing small events may in fact be the best way to prevent large events. But often there are conflicts between preparing for small versus large events. For example, the best response to a small fire in a large building is to stay put until told to move, but at the World Trade Center many learned the hard way that this is bad advice for a large fire. Also, allowing nations to have nuclear weapons can discourage small wars, but encourage large ones.

Similarly, the usual advice for an earthquake is to “duck and cover” under a desk or doorway. This is good advice for small earthquakes, where the main risk is being hit by items falling from the walls or ceiling. But some claim that in a large earthquake where the building collapses, hiding under a desk will most likely get you flattened under that desk; in this case the best place is said to be pressed against the bottom of something incompressible like file cabinets full of paper (Copp, 2000). Unfortunately, our political systems may reward preparing for the most common situations, rather than the greatest expected damage situations.

For some kinds of disruptions, like asteroid strikes, we can work to reduce the rate and severity of events. For other kinds of disruptions, like earthquakes, floods, or hurricanes, we can design our physical systems to better resist damage, such as making buildings that sway rather than crack, and keeping buildings out of flood plains. We can also prevent nuclear proliferation and reduce existing nuclear arsenals.

We can similarly design our social systems to better resist damage. We can consider various crisis situations ahead of time, and make decisions about how to deal with them. We can define who would be in charge of what, and who would have what property rights. We can even create special insurance or crisis management organizations which specialize in dealing with such situations.

If they could count on retaining property rights in a crisis, private organizations would have incentives to set aside private property that they expect to be valuable in such situations. For public goods, or goods with large positive externalities, governments might subsidize

organizations that set aside such goods in preparation for a disaster.

Unfortunately, the fact that large disasters are rare makes it hard to evaluate claims about which mechanisms will actually help in such situations. An engineering organization may claim that a dike would only fail once a century, and police may claim they will keep the peace even with serious social collapse, but track records are not much use in evaluating such claims.

If we value future generations of humanity, we may be willing to take extra efforts to prevent the extinction of humanity. For types of disasters where variations in individual ability to resist disruptions are minor, however, there is little point in explicitly preparing for human extinction possibilities. This is because there is almost no chance that an event of this type would put us very near an extinction borderline. The best we could do here would be to try to prevent all large disruptions. Of course there can be non extinction-related reasons to prepare for such disruptions.

On the other hand, there may be types of disasters where variations in resistance abilities can be important. If so, there might be a substantial chance of finding a post-disaster population that is just above, or just below, a threshold for preserving humanity. In this case it is reasonable to wonder what we might do now to change the odds.

The most obvious possibility would be to create refuges with sufficient resources to help preserve a small group of people through a very large disruption, the resulting social collapse, and a transition period to a post-disaster society. Refuges would have to be strong enough to survive the initial disruption. If desperate people trying to survive a social collapse could threaten a refuge's longterm viability, such as by looting the refuge's resources, then refuges might need to be isolated, well-defended, or secret enough to survive such threats.

We have actually already developed similar refuges to protect social elites during a nuclear war (McCamley, 2007). Though nuclear sanctuaries may not be designed with other human extinction scenarios in mind, it is probably worth considering how they might be adapted to deal with non-nuclear-war disasters. It is also worth considering whether to create a distinct set of refuges, intended for other kinds of disasters. I imagine secret rooms deep in a mine, well stocked with supplies, with some way to monitor the surface and block entry.

A important issue here is whether refuges could by themselves preserve enough humans to supply enough genetic diversity for a post-disaster society. If not, then refuges would either have to count on opening up at the right moment to help preserve enough people outside the sanctuary, or they would need some sort of robust technology for storing genes and implanting them. Perhaps a sperm bank would suffice.

Developing a robust genetic technology might be a challenging task; devices would have to last until the human population reached sufficient size to hold enough genetic diversity on its own. But the payoff could be to drastically reduce the required post-collapse population, perhaps down to a single fertile female. For the purpose of saving humanity reducing the required population from one thousand down to ten is equivalent to a factor of one hundred in current world population, or a factor of one hundred in the severity of each event. In the example of figure 1, it is the same as reducing the disaster event rate by a factor of fifty.

Refuges could in principle hold many kinds of resources which might ease and speed

the restoration of a productive human society. They could preserve libraries, machines, seeds, and much more. But the most important resources would clearly be those that ensure humanity survives. By comparison, on a cosmic scale it is a small matter whether humanity takes one thousand or one hundred thousand years to return to our current level of development. Thus the priority should be resources to support a return to at least a hunter-gatherer society.

It is important to realize that a society rebuilding after a near-extinction crisis would have a vastly smaller scale than our current society; very different types and mixes of capital would be appropriate. Stocking a sanctuary full of the sorts of capital that we find valuable today could be even less useful than the inappropriate medicine, books, or computers often given by first world charities to the third world poor today. Machines would quickly fall into disrepair, and books would impart knowledge that had little practical application.

Instead, one must accept that a very small human population would mostly have to retrace the growth path of our human ancestors; one hundred people cannot support an industrial society today, and perhaps not even a farming society. They might have to start with hunting and gathering, until they could reach a scale where simple farming was feasible. And only when their farming population was large and dense enough could they consider returning to industry.

So it might make sense to stock a refuge with real hunter-gatherers and subsistence farmers, together with the tools they find useful. Of course such people would need to be disciplined enough to wait peacefully in the refuge until the time to emerge was right. Perhaps such people could be rotated periodically from a well protected region where they practiced simple lifestyles, so they could keep their skills fresh. And perhaps we should test our refuge concepts, isolating real people near them for long periods to see how well particular sorts of refuges actually perform at returning their inhabitants to a simple sustainable lifestyle.

## Conclusion

While there are many kinds of catastrophes that might befall humanity, most of the damage that follows large disruptions may come from the ensuing social collapse, rather from the direct effects of the disruption. In thinking about how to prevent and respond to catastrophe, it is therefore crucial to consider the nature of social collapse and how we might minimize it.

After reviewing the nature of society and of social collapse, we have considered how to fit social collapse into a framework where disaster severity follows a reference power law distribution. We made two key distinctions. The first distinction is between types of disasters where small events are the most important, and types of disasters where large events are the most important. The second key distinction is whether individual variation in resistance to a disaster is minor or important.

For types of disaster where both large events and individual resistance variation are important, we have considered some of the tradeoffs involved in trying to preserve humanity. And we have briefly explored the possibility of building special refuges to increase the chances of saving humanity in such situations.

It should go without saying that this has been a very crude and initial analysis; a similar but more careful and numerically precise analysis might be well worth the effort.

## References

- Aghion, P., & Howitt, P. (1998). *Endogenous Growth Theory*. MIT Press, London.
- Barro, R. J., & Sala-I-Martin, X. (2003). *Economic Growth* (2nd edition). MIT Press, London.
- Barton, C., & Nishenko, S. (1997). Natural Disasters: Forecasting Economic and Life Losses.. <http://pubs.usgs.gov/fs/natural-disasters/>.
- Bilham, R. (2004). Urban earthquake fatalities - a safer world or worse to come?. *Seismology Review Letters*.
- Burchell, M. J., Mann, J. R., & Bunch, A. W. (2004). Survival of bacteria and spores under extreme shock pressures. *Monthly Notices of the Royal Astronomical Society*, 352(4), 1273–1278.
- Caplan, B. (2003). The Idea Trap: The Political Economy of Growth Divergence. *European Journal of Political Economy*, 19(2), 183–203.
- Cederman, L.-E. (2003). Modeling the Size of Wars: From Billiard Balls to Sandpiles. *American Political Science Review*, 97(1).
- Christensen, K., Danon, L., Scanlon, T., & Bak, P. (2002). Unified scaling law for earthquakes. *Proceedings of the National Academy of Sciences*, 99(1), 2509–2513.
- Clauset, A., Shalizi, C. R., & Newman, M. E. J. (2007a). Power-law distributions in empirical data.. arXiv:0706.1062v1.
- Clauset, A., Young, M., & Gleditsch, K. S. (2007b). Scale Invariance in the Severity of Terrorism. *Journal of Conflict Resolution*, 5. <http://xxx.lanl.gov/abs/physics/0606007>.
- Copp, D. (2000). Triangle of Life. *American Survival Guide*.
- DeLong, J. B. (2006). Growth is Good. *Harvard Magazine*, 19–20.
- deMenocal, P. B. (2001). Cultural Responses to Climate Change During the Late Holocene. *Science*, 292(5517), 667 – 673.
- Diamond, J. (2005). *Collapse: How Societies Choose to Fail or Succeed*. Viking Adult, New York.
- Grubler, A. (1998). *Technology and Global Change*. Cambridge University Press, New York.

- Hanson, R. (2000). Long-term growth as a sequence of exponential modes..
- Haug, G. H., Gnther, D., Peterson, L. C., Sigman, D. M., Hughen, K. A., & Aeschlimann, B. (2003). Climate and the Collapse of Maya Civilization. *Science*, *299*(5613), 1731 – 1735.
- Hey, J. (2005). On the Number of New World Founders: A Population Genetic Portrait of the Peopling of the Americas. *PLoS Biology*, *3*(6), 965–975.
- Jones, C. I. (2002). *Introduction to Economic Growth* (2nd edition). W. W. Norton & Company.
- Lay, T., & Wallace, T. (1995). *Modern Global Seismology*. Academic Press, San Deigo, California.
- McCamley, N. (2007). *Cold War Secret Nuclear Bunkers*. Pen and Sword.
- Morrison, D., Harris, A. W., Sommer, G., Chapman, C. R., & Carusi, A. (2003). Dealing with the Impact Hazard. In Bottke, W., Cellino, A., Paolicchi, P., & Binzel, R. P. (Eds.), *Asteroids III*. University of Arizona Press, Tuscon.
- Murray-McIntosh, R. P., Scrimshaw, B. J., Hatfield, P. J., & Penny, D. (1998). Testing migration patterns and estimating founding population size in Polynesia by using human mtDNA sequences. *Proceedings of National Academy of Science USA*, *95*, 90479052.
- Nishenko, S., & Barton, C. (1995). Scaling Laws for Natural Disaster Fatalities. In Rundle, J., Klein, F., & Turcotte, D. (Eds.), *Reduction and Predictability of Natural Disasters*, Vol. 25, p. 32. Addison Wesley.
- Posner, R. A. (2004). *Catastrophe : risk and response*. Oxford University Press, New York.
- Rhodes, C. J., Jensen, H. J., & Anderson, R. M. (1997). On the Critical Behaviour of Simple Epidemics. *Proceedings of the Royal Society B: Biological Sciences*, *264*(1388), 1639–1646.
- Rodrik, D. (1999). Where Did All the Growth Go? External Shocks, Social Conflict, and Growth Collapses. *Journal of Economic Growth*, *4*(4), 385–412.
- Sanders, D. E. A. (2005). The Modeling of Extreme Events..
- Tainter, J. (1988). *The Collapse of Complex Societies*. Cambridge University Press, New York.
- Turcotte, D. L. (1999). Self-organized criticality. *Reports on Progress in Physics*, *62*, 13771429.

Watts, D., Muhamad, R., Medina, D., & Dodds, P. (2005). Multiscale, resurgent epidemics in a hierarchical metapopulation model. *Proceedings of the National Academy of Sciences*, *102*(32), 11157–11162.

Weiss, H., & Bradley, R. S. (2001). What Drives Societal Collapse?. *Science*, *291*(5504), 609–610.