

# Book Review: Ages Of Discord

*Posted on September 2, 2019 by Scott Alexander*




## I


I recently reviewed [Secular Cycles](#), which presents a demographic-structural theory of the growth and decline of pre-industrial civilizations. When land is plentiful, population grows and the economy prospers. When land reaches its carrying capacity and income declines to subsistence, the area is at risk of famines, diseases, and wars – which kill enough people that land becomes plentiful again. During good times, elites prosper and act in unity; during bad times, elites turn on each other in an age of backstabbing and civil strife. It seemed pretty reasonable, and authors Peter Turchin and Sergey Nefedov had lots of data to support it. *Ages of Discord* is Turchin's attempt to apply the same theory to modern America. There are many reasons to think this shouldn't work, and the book does a bad job addressing them. So I want to start by presenting Turchin's data showing such cycles exist, so we can at least see why the hypothesis might be tempting. Once we've seen the data, we can decide how turned off we want to be by the theoretical problems.

The first of Turchin's two cyclic patterns is a long cycle of national growth and decline. In *Secular Cycles'* pre-industrial societies, this

pattern lasted about 300 years; in *Ages of Discord's* picture of the modern US, it lasts about 150:




This summary figure combines many more specific datasets. For example, archaeologists frequently assess the prosperity of a period by the heights of its skeletons. Well-nourished, happy children tend to grow taller; a layer with tall skeletons probably represents good times during the relevant archaeological period; one with stunted skeletons probably represents famine and stress. What if we applied this to the modern US?



*Average US height and life expectancy over time.*


As far as I can tell, the height graph is raw data. The life expectancy graph is the raw data minus an assumed constant positive trend – that is, given that technological advance is increasing life expectancy at a linear rate, what are the other factors you see when you subtract that out? The exact statistical logic be buried in Turchin's source (**Historical Statistics of the United States**, Carter et al 2004), which I don't have and can't judge.

This next graph is the median wage divided by GDP per capita, a crude measure of income equality:




**Lower** values represent **more** inequality.


This next graph is median female age at first marriage. Turchin draws on research suggesting this tracks social optimism. In good times, young people can easily become independent and start supporting a family; in bad times, they will want to wait to make sure their lives are stable before settling down:




This next graph is Yale tuition as a multiple of average manufacturing worker income. To some degree this will track inequality in general, but Turchin thinks it also measures something like “difficulty of upward mobility”:



This next graph shows DW-NOMINATE's "Political Polarization Index", a complicated metric occasionally used by historians of politics. It measures the difference in voting patterns between the average Democrat in Congress and the average Republican in Congress (or for periods before the Democrats and Republicans, whichever two major parties there were). During times of low partisanship, congressional votes will be dominated by local or individual factors; during times of high partisanship, it will be dominated by party identification:



I've included only those graphs which cover the entire 1780 – present period; the book includes many others that only cover shorter intervals (mostly the more recent periods when we have better data). All of them, including the shorter ones not included here, reflect the same general pattern. You can see it most easily if you standardize all the indicators to the same scale, match the signs so that up always means good and down always means bad, and put them all together:




*Note that these aren't exactly the same indicators I featured above;  
we'll discuss immigration later.*

The “average” line on this graph is the one that went into making the summary graphic above. Turchin believes that after the American Revolution, there was a period of instability lasting a few decades (eg Shays’ Rebellion, Whiskey Rebellion) but that America reached a maximum of unity, prosperity, and equality around 1820. Things gradually got worse from there, culminating in a peak of inequality, misery, and division around 1900. The reforms of the Progressive Era gradually made things better, with another unity/prosperity/equality maximum around 1960. Since then, an increasing confluence of negative factors (named here as the Reagan Era trend reversal, but Turchin admits it began before Reagan) has been making things worse again.



## II


Along with this “grand cycle” of 150 years, Turchin adds a shorter instability cycle of 40-60 years. This is the same 40-60 year instability cycle that appeared in *Secular Cycles*, where Turchin called it “the bigenerational cycle”, or the “fathers and sons cycle”.



*Timing and intensity of internal war in medieval and early modern England, from Turchin and Nefedov 2009.*


The derivation of this cycle, explained on pages 45 – 58 of *Ages of Discord*, is one of the highlights of the book. Turchin draws on the kind of models epidemiologists use to track pandemics, thinking of violence as an infection and radicals as plague-bearers. You start with an unexposed vulnerable population. Some radical – patient

zero – starts calling for violence. His ideas spread to a certain percent of people he interacts with, gradually “infecting” more and more people with the “radical ideas” virus. But after enough time radicalized, some people “recover” – they become exhausted with or disillusioned by conflict, and become pro-cooperation “active moderates” who are impossible to reinfect (in the epidemic model, they are “inoculated”, but they also have an ability without a clear epidemiological equivalent to dampen radicalism in people around them). As the rates of radicals, active moderates, and unexposed dynamically vary, you get a cyclic pattern. First everyone is unexposed. Then radicalism gradually spreads. Then active moderation gradually spreads, until it reaches a tipping point where it triumphs and radicalism is suppressed to a few isolated reservoirs in the population. Then the active moderates gradually die off, new unexposed people are gradually born, and the cycle starts again. Fiddling with all these various parameters, Turchin is able to get the system to produce 40-60 year waves of instability.



To check this empirically, Turchin tries to measure the number of “instability events” in the US over various periods. He very correctly tries to use lists made by others (since they are harder to bias), but when people haven’t catalogued exactly the kind of instability he’s interested in over the entire 1780 – present period, he sometimes adds his own interpretation. He ends up summing riots, lynchings, terrorism (including assassinations), and mass shootings – you can see his definition for each of these starting on page 114; the short version is that all the definitions seem reasonable but inevitably include a lot of degrees of freedom.

When he adds all this together, here’s what happens:



*Political instability / violent events show three peaks, around 1870, 1920, and 1970.*

The 1870 peak includes the Civil War, various Civil War associated violence (eg draft riots), and the violence around Reconstruction (including the rise of the Ku Klux Klan and related violence to try to control newly emancipated blacks).

The 1920 peak includes the height of the early US labor movement. Turchin discusses the Mine War, an “undeclared war” from 1920-1921 between bosses and laborers in Appalachian coal country:

Although it started as a labor dispute, it eventually turned into the largest armed insurrection in US history, other than the Civil War. Between 10,000 and 15,000 miners armed with rifles fought thousands of strike-breakers and sheriff's deputies, called the Logan Defenders. The insurrection was ended by the US Army. While such violent incidents were exceptional, they took place against a background of a general "class war" that had been intensifying since the violent teens. "In 1919 nearly four million workers (21% of the workforce) took disruptive action in the face of employer reluctance to recognize or bargain with unions" (Domhoff and Webber, 2011:74).

Along with labor violence, 1920 was also a peak in racial violence:

Race-motivated riots also peaked around 1920. The two most serious such outbreaks were the Red Summer of 1919 (McWhirter 2011) and the Tulsa (Oklahoma) Race Riot. The Red Summer involved riots in more than 20 cities across the United States and resulted in something like 1,000 fatalities. The Tulsa riot in 1921, which caused about 300 deaths, took on an aspect of civil war, in which thousands of whites and blacks, armed with firearms, fought in the streets, and most of the Greenwood District, a prosperous black neighborhood, was destroyed.

And terrorism:

The bombing campaign by Italian anarchists (“Galleanists”) culminated in the 1920 explosion on Wall Street, which caused 38 fatalities.

The same problems: labor unrest, racial violence, terrorism – repeated during the 1970s spike. Instead of quoting Turchin on this, I want to quote [this Status 451 review](#) of *Days of Rage*, because it blew my mind:

“People have completely forgotten that in 1972 we had over nineteen hundred domestic bombings in the United States.”  
— Max Noel, FBI (ret.)

Recently, I had my head torn off by a book: Bryan Burrough’s *Days of Rage*, about the 1970s underground. It’s the most important book I’ve read in a year. So I did a series of running tweetstorms about it, and Clark asked me if he could collect them for posterity. I’ve edited them slightly for editorial coherence.

*Days of Rage* is important, because this stuff is forgotten and it shouldn’t be. The 1970s underground wasn’t small. It was hundreds of people becoming urban guerrillas. Bombing buildings: the Pentagon, the Capitol, courthouses, restaurants, corporations. Robbing banks. Assassinating police. People really thought that revolution was imminent, and thought violence would bring it about.

One thing that Burrough returns to in Days of Rage, over and over and over, is how forgotten so much of this stuff is. Puerto Rican separatists bombed NYC like 300 times, killed people, shot up Congress, tried to kill POTUS (Truman). Nobody remembers it.

The passage speaks to me because – yeah, nobody remembers it. This is also how I feel about the 1920 spike in violence. I'd heard about the Tulsa race riot, but the Mine War and the bombing of Wall Street and all the other stuff was new to me. This matters because my intuitions before reading this book would not have been that there were three giant spikes in violence/instability in US history located fifty years apart. I think the lesson I learn is not to trust my intuitions, and to be a little more sympathetic to Turchin's data.

One more thing: the 1770 spike was obviously the American Revolution and all of the riots and communal violence associated with it (eg against Tories). Where was the 1820 spike? Turchin admits it didn't happen. He says that because 1820 was the absolute best part of the 150 year grand cycle, everybody was so happy and well-off and patriotic that the scheduled instability peak just fizzled out. Although Turchin doesn't mention it, you could make a similar argument that the 1870 spike was especially bad (see: the entire frickin' Civil War) because it hit close to (though not exactly at) the worst part of the grand cycle. 1920 hit around the middle, and 1970 during a somewhat-good period, so they fell in between the nonissue of 1820 and the disaster of 1870.

### III

I haven't forgotten the original question – what drives these 150 year cycles of rise and decline – but I want to stay with the data just a little longer. Again, these data are really interesting. Either some sort of really interesting theory has to be behind them – or they're just low-quality data cherry-picked to make a point. Which are they? Here are a couple of [spot-checks](#) to see if the data are any good.

**First spot check:** can I confirm Turchin's data from independent sources?

- [Here](#) is a graph of average US height over time which seems broadly similar to Turchin's.
- [Here](#) is a different measure of US income inequality over time, which again seems broadly similar to Turchin's. [Piketty](#) also presents very similar data, though his story places more emphasis on the World Wars and less on the labor movement.
- The Columbia Law Review measures [political polarization over time](#) and gets mostly the same numbers as Turchin.

I'm going to consider this successfully checked; Turchin's data all seem basically accurate.



**Second spot check:** do other indicators Turchin didn't include confirm the pattern he detects, or did he just cherry-pick the data series that worked? Spoiler: I wasn't able to do this one. It was too hard to think of measures that should reflect general well-being and that we have 200+ years of unconfounded data for. But here are my various failures:

- The [annual improvement in mortality rate](#) does not seem to follow the cyclic pattern. But isn't this more driven by a few random factors like smoking rates and the logic of technological advance?
- [Treasury bonds](#) maybe kind of follow the pattern until 1980, after which they go crazy.
- [Divorce rates](#) look kind of iffy, but isn't that just a bunch of random factors?
- [Homicide rates](#), with the general downward trend removed, sort of follow the pattern, except for the recent decline?
- [USD/GBP exchange rates](#) don't show the pattern at all, but that could be because of things going on in Britain?

The thing is – really I have no reason to expect divorce rates, homicide rates, exchange rates etc to track national flourishing. For one thing, they may just be totally unrelated. For another, even if they were tenuously related, there are all sorts of other random factors that can affect them. The problem is, I would have said this was true for height, age at first marriage, and income inequality too, be-

fore Turchin gave me convincing-sounding stories for why it wasn't. I think my lesson is that I have no idea which indicators should vs. shouldn't follow a secular-cyclic pattern and so I can't do this spot check against cherry-picking the way I hoped.

**Third spot check:** common sense. Here are some things that stood out to me:

- The Civil War is at a low-ish part of the cycle, but by no means the lowest.
- The Great Depression happened at a medium part of the cycle, when things should have been quickly getting better.
- Even though there was a lot of new optimism with Reagan, continuing through the Clinton years, the cycle does not reflect this at all.

Maybe we can rescue the first and third problem by combining the 150 year cycle with the shorter 50 year cycle. The Civil War was determined by the 50-year cycle having its occasional burst of violence at the same time the 150-year cycle was at a low-ish point. People have good memories of Reagan because the chaos of the 1970 violence burst had ended.

As for the second, Turchin is aware of the problem. He writes:

There is a widely held belief among economists and other social scientists that the 1930s were the “defining moment”

in the development of the American politico-economic system (Bordo et al 1998). When we look at the major structural-demographic variables, however, the decade of the 1930s does not seem to be a turning point. Structural-demographic trends that were established during the Progressive Era continues through the 1930s, although some of them accelerated.

Most notably, all the well-being variables that went through trend reversals before the Great Depression – between 1900 and 1920. From roughly 1910 and to 1960 they all increased roughly monotonically, with only one or two minor fluctuations around the upward trend. The dynamics of real wages also do not exhibit a breaking point in the 1930s, although there was a minor acceleration after 1932.

By comparison, he plays up the conveniently-timed (and hitherto unknown to me) depression of the mid-1890s. Quoting Turchin quoting McCormick:

No depression had ever been as deep and tragic as the one that lasted from 1893 to 1897. Millions suffered unemployment, especially during the winters of 1893-4 and 1894-5, and thousands of ‘tramps’ wandered the countryside in search of food [...]

Despite real hardship resulting from massive unemployment, well-being indicators suggest that the human cost of the Great Depression of the 1930s did not match that of the

“First Great Depression” of the 1890s (see also Grant 1983:3-11 for a general discussion of the severity of the 1890s depression. Furthermore, while the 1930s are remembered as a period of violent labor unrest, the intensity of class struggle was actually lower than during the 1890s depression. According to the US Political Violence Database (Turchin et al. 2012) there were 32 lethal labor disputes during the 1890s that collectively caused 140 deaths, compared with 20 such disputes in the 1930s with the total of 55 deaths. Furthermore, the last lethal strike in US labor history was in 1937... in other words, the 1930s was actually the last uptick of violent class struggle in the US, superimposed on an overall declining trend.

The 1930s Depression is probably remembered (or rather misremembered) as the worst economic slump in US history, simply because it was the last of the great depressions of the post-Civil War era.

**Fourth spot check:** Did I randomly notice any egregious errors while reading the book?

On page 70, Turchin discusses “the great cholera epidemic of 1849, which carried away up to 10% of the American population”. This seemed unbelievably high to me. I checked the source he cited, Kohl’s “Encyclopedia Of Plague And Pestilence”, which did give that number. But every other source I checked agreed that the epidemic “only” killed between 0.3% – 1% of the US population (it did hit 10% in a few especially unlucky cities like St. Louis). I cannot

fault Turchin's scholarship in the sense of correctly repeating something written in an encyclopedia, but unless I'm missing something I do fault his common sense.

Also, on page 234, Turchin interprets the percent of medical school graduates who get a residency as "the gap between the demand and supply of MD positions", which he ties into a wider argument about elite overproduction. But I think this shows a limited understanding of how the medical system works. There is currently a severe *undersupply* of doctors – try getting an appointment with a specialist who takes insurance in a reasonable amount of time if you don't believe me. Residencies aren't limited by organic demand. They're limited because the government places so many restrictions on them that hospitals don't sponsor them without government funding, and the government is too stingy to fund more of them. None of this has anything to do with elite overproduction.

These are just two small errors in a long book. But they're two errors in medicine, the field I know something about. This makes me worry about [Gell-Mann Amnesia](#): if I notice errors in my own field, how many errors must there be in other fields that I just didn't catch?

My overall conclusion from the spot-checks is that the data as presented are basically accurate, but that everything else is so dependent on litigating which things are vs. aren't in accordance with the theory that I basically give up.

## IV

Okay. We've gone through the data supporting the grand cycle. We've gone through the data and theory for the 40-60 year instability cycle. We've gone through the reasons to trust vs. distrust the data. Time to go back to the question we started with: why should the grand cycle, originally derived from the Malthusian principles that govern pre-industrial societies, hold in the modern US? Food and land are no longer limiting resources; famines, disease, and wars no longer substantially decrease population. Almost every factor that drives the original secular cycle is missing; why even consider the possibility that it might still apply?

I've put this off because, even though this is the obvious question *Ages of Discord* faces from page one, I found it hard to get a single clear answer.

Sometimes, Turchin talks about the supply vs. demand of labor. In times when the supply of labor outpaces demand, wages go down, inequality increases, elites fragment, and the country gets worse, mimicking the "land is at carrying capacity" stage of the Malthusian cycle. In times when demand for labor exceeds supply, wages go up, inequality decreases, elites unite, and the country gets better. The government is controlled by plutocrats, who always want wages to be low. So they implement policies that increase the supply of labor, especially loose immigration laws. But their actions cause inequality to increase and everyone to become miserable. Ordinary people organize resistance: populist movements, socialist cadres, labor unions. The system teeters on the edge of violence,

revolution, and total disintegration. Since the elites don't want those things, they take a step back, realize they're killing the goose that lays the golden egg, and decide to loosen their grip on the neck of the populace. The government becomes moderately pro-labor and progressive for a while, and tightens immigration laws. The oversupply of labor decreases, wages go up, inequality goes down, and everyone is happy. After everyone has been happy for a while, the populists/socialists/unions lose relevance and drift apart. A new generation of elites who have never felt threatened come to power, and they think to themselves "What if we used our control of the government to squeeze labor harder?" Thus the cycle begins again.

But at other times, Turchin talks more about "elite overproduction". When there are relatively few elites, they can cooperate for their common good. Bipartisanship is high, everyone is unified behind a system perceived as wise and benevolent, and we get a historical period like the 1820s US golden age that historians call [The Era Of Good Feelings](#). But as the number of elites outstrips the number of high-status positions, competition heats up. Elites realize they can get a leg up in an increasingly difficult rat race by backstabbing against each other and the country. Government and culture enter a defect-defect era of hyperpartisanship, where everyone burns the commons of productive norms and institutions in order to get ahead. Eventually... some process reverses this or something?... and then the cycle starts again.


At still other times, Turchin seems to retreat to a sort of mathematical formalism. He constructs an extremely hokey-looking dy-

dynamic feedback model, based on ideas like “assume that the level of discontent among ordinary people equals the urbanization rate x the age structure x the inverse of their wages relative to the elite” or “let us define the fiscal distress index as  $\text{debt} \div \text{GDP} \times \text{the level of distrust in state institutions}$ ”. Then he puts these all together into a model that calculates how the the level of discontent affects and is affected by the level of state fiscal distress and a few dozen other variables. On the one hand, this is really cool, and watching it in action gives you the same kind of feeling Seldon must have had inventing psychohistory. On the other, it seems *really* made-up. Turchin admits that dynamic feedback systems are infamous for going completely haywire if they are even a tiny bit skew to reality, but assures us that he understands the cutting-edge of the field and how to make them not to do that. I don’t know enough to judge whether he’s right or wrong, but my priors are on “extremely, almost unfathomably wrong”. Still, at times he reminds us that the shifts of dynamic feedback systems can be attributed only to the system in its entirety, and that trying to tell stories about or point to specific factors involved in any particular shift is an approximation at best.

All of these three stories run into problems almost immediately.

**First**, the supply of labor story focuses pretty heavily on immigration. Turchin puts a lot of work into showing that immigration follows the secular cycle patterns; it is highest at the worst part of the cycle, and lowest at the best parts:





In this model, immigration is a tool of the plutocracy. High supply of labor (relative to demand) drives down wages, increases inequality, and lowers workers' bargaining power. If the labor supply is poorly organized, comes from places that don't understand the concept of "union", don't know their rights, and have racial and linguistic barriers preventing them from cooperating with the rest of the working class, well, even better. Thus, periods when the plutocracy is successfully squeezing the working class are marked by high immigration. Periods when the plutocracy fears the working class and feels compelled to be nice to them are marked by low immigration.

This position makes some sense and is loosely supported by the long-term data above. But isn't this one of the most-studied topics in the history of economics? Hasn't it been proven almost beyond doubt that immigrants don't steal jobs from American workers, and that since they consume products themselves (and thus increase

the demand for labor) they don't affect the supply/demand balance that sets wages?

It appears I might just be totally miscalibrated on this topic. I checked the [IGM Economic Experts Panel](#). Although most of the expert economists surveyed believed immigration was a net good for America, they did say (50% agree to only 9% disagree) that “unless they were compensated by others, many low-skilled American workers would be substantially worse off if a larger number of low-skilled foreign workers were legally allowed to enter the US each year”. I'm having trouble seeing the difference between this statement (which economists seem very convinced is true) and “you should worry about immigrants stealing your job” (which everyone seems very convinced is false). It might be something like – immigration generally makes “the economy better”, but there's no guarantee that these gains are evenly distributed, and so it can be bad for low-skilled workers in particular? I don't know, this would still represent a pretty big update, but given that I was told all top economists think one thing, and now I have a survey of all top economists saying the other, I guess big updates are unavoidable. Interested in hearing from someone who knows more about this.

Even if it's true that immigration can hurt low-skilled workers, Turchin's position – which is that increased immigration is responsible for a very large portion of [post-1973 wage stagnation](#) and the recent trend toward rising inequality – sounds shocking to current political sensibilities. But all Turchin has to say is:

An imbalance between labor supply and demand clearly played an important role in driving real wages down after 1978. As Harvard economist George J. Borjas recently wrote, “The best empirical research that tries to examine what has actually happened in the US labor market aligns well with economic theory: An increase in the number of workers leads to lower wages.”

My impression was that Borjas was an increasingly isolated contrarian voice, so once again, I just don’t know what to do here.

**Second**, the plutocratic oppression story relies pretty heavily on the idea that inequality is a unique bad. This fits the zeitgeist pretty well, but it’s a little confusing. Why should commoners care about their wages relative to elites, as opposed to their absolute wages? Although median-wage-relative-to-GDP has gone down over the past few decades, absolute median wage has gone up – just a little, slowly enough that it’s rightly considered a problem – but it has gone up. Since modern wages are well above 1950s wages, in what sense should modern people feel like they are economically bad off in a way 1950s people didn’t? This isn’t a problem for Turchin’s theory so much as a general mystery, but it’s a general mystery I care about a lot. One answer is that the [cost disease](#) is fueled by a Baumol effect pegged to per capital income (see part 3 [here](#)), and this is a way that increasing elite wealth can absolutely (not relatively) immiserate the lower classes.

Likewise, what about [The Spirit Level Delusion](#) and other resources showing that, across countries, inequality is not particularly corre-

lated with social bads? Does this challenge Turchin's America-centric findings that everything gets worse along with inequality levels?

**Third**, the plutocratic oppression story meshes poorly with the elite overproduction story. In elite overproduction, united elites are a sign of good times to come; divided elites means dysfunctional government and potential violence. But as [Pseudoerasmus points out](#), united elites are often united *against* the commoners, and we should expect inequality to be highest at times when the elites are able to work together to fight for a larger share of the pie. But I think this is the opposite of Turchin's story, where elites unite only to make concessions, and elite unity equals popular prosperity.

**Fourth**, everything about the elite overproduction story confuses me. Who are "elites"? This category made sense in *Secular Cycles*, which discussed agrarian societies with a distinct titled nobility. But Turchin wants to define US elites in terms of wealth, which follows a continuous distribution. And if you're defining elites by wealth, it doesn't make sense to talk about "not enough high-status positions for all elites"; if you're elite (by virtue of your great wealth), by definition you already have what you need to maintain your elite status. Turchin seems aware of this issue, and sometimes talks about "elite aspirants" – some kind of upper class who expect to be wealthy, but might or might not get that aspiration fulfilled. But then understanding elite overproduction hinges on what makes one non-rich-person person a commoner vs. another non-rich-person an "elite aspirant", and I don't remember any clear discussion of this in the book.

**Fifth**, what drives elite overproduction? Why do elites (as a percent of the population) increase during some periods and decrease during others? Why should this be a cycle rather than a random walk?

My guess is that *Ages of Discord* contains answers to some of these questions and I just missed them. But I missed them after reading the book pretty closely to try to find them, and I didn't feel like there were any similar holes in *Secular Cycles*. As a result, although the book had some fascinating data, I felt like it lacked a clear and lucid thesis about exactly what was going on.

## V

Accepting the data as basically right, do we have to try to wring some sense out of the theory?

The data cover a cycle and a half. That means we only sort of barely get to see the cycle "repeat". The conclusion that it is a cycle and not some disconnected trends is based only on the single coincidence that it was 70ish years from the first turning point (1820) to the second (1890), and also 70ish years from the second to the third (1960).

A parsimonious explanation would be "for some reason things were going unusually well around 1820, unusually badly around 1890, and unusually well around 1960 again." This is actually really interesting – I didn't know it was true before reading this book,

and it changes my conception of American history a lot. But it's a lot less interesting than the discovery of a secular cycle.

I think the parsimonious explanation is close to what Thomas Piketty argued in his [\*Capital In The Twenty-First Century\*](#). Inequality was rising until the World Wars, because that's what inequality naturally does given reasonable assumptions about growth rates. Then the Depression and World Wars wiped out a lot of existing money and power structures and made things equal again for a little while. Then inequality started rising again, because that's what inequality naturally does given reasonable assumptions about growth rates. Add in a pinch of *The Spirit Level* – inequality is a mysterious magic poison that somehow makes everything else worse – and there's not much left to be explained.

(some exceptions: why was inequality decreasing until 1820? Does inequality *really* drive political polarization? When immigration corresponds to periods of high inequality, is the immigration *causing* the inequality? And what about the 50 year cycle of violence? That's another coincidence we didn't include in the coincidence list!)

So what can we get from *Ages of Discord* that we can't get from Piketty?

First, the concept of “elite overproduction” is one that worms its way into your head. It's the sort of thing that was constantly in the background of [\*Increasingly Competitive College Admissions: Much More Than You Wanted To Know\*](#). It's the sort of thing you think

about when a million fresh-faced college graduates want to become Journalists and Shape The Conversation and Fight For Justice and realistically just end up getting ground up and spit out by clickbait websites. *Ages of Discord* didn't do a great job breaking down its exact dynamics, but I'm grateful for its work bringing it from a sort of shared unconscious assumption into the light where we can talk about it.

Second, the idea of a deep link between various indicators of goodness and badness – like wages and partisan polarization – is an important one. It forces me to reevaluate things I had considered settled, like that immigration doesn't worsen inequality, or that inequality is not a magical curse that poisons everything.

Third, historians have to choose what events to focus on. Normal historians usually focus on the same normal events. Unusual historians sometimes focus on neglected events that support their unusual theses, so reading someone like Turchin is a good way to learn parts of history you'd never encounter otherwise. Some of these I was able to mention above – like the Mine War of 1920 or the cholera epidemic of 1849; I might make another post for some of the others.

Fourth, it tries to link events most people would consider separate – wage stagnation since 1973, the Great Stagnation in technology, the decline of Peter Thiel's ["definite optimism"](#), the rise of partisan polarization. I'm not sure exactly *how* it links them or what it has to say about the link, but link them it does.

But the most important thing about this book is that Turchin claims to be able to predict the future. The book (written just before Trump was elected in 2016) ends by saying that “we live in times of intensifying structural-demographic pressures for instability”. The next bigenerational burst of violence is scheduled for about 2020 (realistically +/- a few years). It’s at a low point in the grand cycle, so it should be a doozy.

What about beyond that? It’s unclear *exactly* where he thinks we are right now in the grand cycle. If the current cycle lasts exactly as long as the last one, we would expect it to bottom out in 2030, but Turchin never claims every cycle is exactly as long. A few of his graphs suggest a hint of curvature, suggesting we might currently be in the worst of it. The socialists seem to have gotten their act together and become an important political force, which the theory predicts is a necessary precursor to change.

I think we can count the book as having made correct predictions if violence spikes in the very near future (are the current number of mass shootings enough to satisfy this requirement? I would have to see it graphed using the same measurements as past spikes), and if sometime in the next decade or so things start looking like there’s a ray of light at the end of the tunnel.

I am pretty interested in finding other ways to test Turchin’s theories. I’m going to ask some of my math genius friends to see if the dynamic feedback models check out; if anyone wants to help, let me know how I can help you (if money is an issue, I can send you a copy of the book, and I will definitely publish anything you find on



this blog). If anyone has any other ideas for to indicators that should be correlated with the secular cycle, and ideas about how to find them, I'm intereted in that too. And if anyone thinks they can explain the elite overproduction issue, please enlighten me.

I ended my review of *Secular Cycles* by saying:

One thing that strikes me about [Turchin]'s cycles is the ideological component. They describe how, during a growth phase, everyone is optimistic and patriotic, secure in the knowledge that there is enough for everybody. During the stagflation phase, inequality increases, but concern about inequality increases even more, zero-sum thinking predominates, and social trust craters (both because people are actually defecting, and because it's in lots of people's interest to play up the degree to which people are defecting). By the crisis phase, partisanship is much stronger than patriotism and radicals are talking openly about how violence is ethically obligatory.

And then, eventually, things get better. There is a new Augustan Age of virtue and the reestablishment of all good things. This is a really interesting claim. Western philosophy tends to think in terms of trends, not cycles. We see everything going on around us, and we think this is some endless trend towards more partisanship, more inequality, more hatred, and more state dysfunction. But *Secular Cycles* offers a narrative where endless trends can end, and things can get better after all.

This is still the hope, I guess. I don't have a lot of faith in human effort to restore [niceness, community, and civilization](#). All I can do is pray the Vast Formless Things accomplish it for us without asking us first.