G-FEED

Global Food, Environment and Economic Dynamics

Who we are

We are an interdisciplinary group working cooperatively to understand the relationship between society and the environment.

We examine regional and global scale phenomena using statistical analyses of real world data.

More about us.

Contributors

Maximilian Auffhammer
Marshall Burke
Solomon Hsiang
David Lobell
Michael Roberts
Wolfram Schlenker

Search This Blog

Search

Links

G-FEED.org
Greed, Green & Grains
Fight Entropy
Wolfram Schlenker's Blog

Subscribe by RSS

Posts
Comments

Subscribe by Email

Email address... Submit

Translate G-FEED

Select Language

Labels

- climate change (36)
- agriculture (35)
- empirical research (30)
- research summary (25)
- climate (24)
- adaptation (14)
- yields (14)
- conflict (13)
- drought (13)
- statistics (13)
- media (12)
- heat (10)
- economics (9)
- corn (8)
- extreme heat (7)
- forecasts (6)
- methods (6)

Monday, October 26, 2015

Climate change and the global economy

[image lifted from PBS press coverage]

Sol, Ted Miguel, and I have a new paper out in *Nature* that looks at at how past and future temperature changes might affect global economic output. In particular, we study the historical relationship between temperature and country-level output with an eye for potential non-linearities in the macro data (which have cropped up everywhere in the micro literature). We then combine historical results with global climate model estimates of future warming to come up with some projections of the potential future impacts of warming. We wrap up by trying to compare our damage estimates to the damage functions currently in the Integrated Assessment Models (IAMs).

We get some big numbers. Looking historically, we see that output in both rich and poor countries alike has been shaped by changes in temperature, and that temperature appears to affect growth rate of per capita GDP and not just the level of GDP (which matters a whole lot when you do the projections). Importantly, we don't see big differences between rich and poor countries in how they respond to changes in temperature historically. Differences we do see across countries appear driven more by countries' average temperatures than by their average incomes, with cooler countries growing faster on average during yeas that are warmer-than average for them, and hotter countries growing slower.

This non-linear response peaks at an annual average temperature of around 13C which just so happens to be the annual average temperature of both Palo Alto and New York City. (For the naysayers: There is nothing mechanical in this fact; we can drop the US from the country-level regressions and we get the same optimum of 13C). The effect of temperature on growth rates is pretty flat around this peak, but gets pretty steep as you move away from the optimum in either direction. We find that for really hot or really cold countries, +1C changes in annual temperature have historically moved growth rates up (for cold countries) or down (for hot countries) by a percentage point -- i.e. a hot country goes from growing at 2% per year to 1% a year. That is a big number. And for poor countries, it's basically what was shown in the seminal Dell, Jones, Olken piece from a few years ago. The big difference is we see a lot more action in the richer, cooler countries than DJO found in their earlier paper -- a difference we spend a lot of time exploring in the supplement to our paper.

We then run the world forward under a RCP8.5, a business-as-usual emissions scenario that makes the world pretty hot by end of century (+4.3C is the population-weighted projected temperature increase under RCP8.5 by 2100 that we pull off the models). Doing this requires three pieces of data, as we describe here [after clicking on a country, scroll down to the "how do we arrive at these numbers" section]. Cool high-latitude countries warm up a bunch (much more than global average) and so could stand to benefit substantially from climate change -- recall the +1%/C marginal effect from above. Countries at or beyond the optimum are harmed, and increasingly so as the temperature heats up. Globally, we find that under RCP8.5, the global economy could be more than 20% smaller by 2100 than it would have been had temperatures remained fixed at today's values. This does not mean that the world will be poorer in 2100. It almost certainly will not. It means that it will be less rich than it would have been had temperatures not warmed.

Importantly, our estimates mechanically do not include the potential future effects of stuff we have not observed historically (e.g. sea level rise), nor do they contain non-marketed things we care about that do not show up in GDP (e.g. polar bears).

We built a little website to take people through the paper and let people play with the country-level results. We've posted all our data and replication code, and would love people to show us where we went wrong.

In the spirit of earlier blog posts, we again want to use this space to respond to some of the early comments and criticism we've gotten on the paper. This accompanies a related attempt to answer some "frequently asked questions" about our paper that we got from the press and earlier inquirires. We imagine that we will be updating this blog post as additional criticism rolls in. But here are some of the "Frequently Heard Criticisms" (FHCs) so far, and some responses:

- news (6)
- prices (6)
- temperature (6)
- commodity (5)
- rainfall (5)
- degree days (4)
- nonlinear temperature effects (4)
- uncertainty (4)
- usda (4)
- variance (4)
- commodity prices (3)
- development (3)
- externalities (3)
- extreme events (3)
- remote sensing (3)
- speculation (3)
- world supply (3)
- yield gap (3)
- ENSO (2)
- Obama (2)
- cell towers (2)
- daily weather data (2)
- ethanol (2)
- guest post (2)
- hunger (2)
- insurance (2)
- irrigation (2)
- mitigation (2)
- nitrogen (2)
- policy (2)
- satellites (2)
- storage. (2)
- water quality (2)
- world demand (2)
- Angry Birds (1)
- CAFO (1)
- Dr. Seuss (1)
- FAO (1)
- Facebook (1)
- Genghis Khan (1)
- Google (1)
- PRISM (1)
- air pollution (1)
- auctions (1)
- bubble (1)
- conservation (1)
- crop disease (1)
- crop rotations (1)
- data (1)
- demand (1)
- double crop (1)
- endogeneity (1)
- fertilizer (1)
- fixed effects (1)
- growth (1)
- hadoop (1)
- hurricanes (1)
- illusions (1)
- india (1)
- intercomparisons (1)
- ipcc (1)
- land values (1)
- nice suit (1)
- ozone (1)
- payments for environmental services (1)
- phosphorus (1)
- plastics (1)
- poetry (1)
- poop (1)
- public finance (1)
- retail (1)
- scientific consensus (1)

1. **These results just don't pass the "sniff test" [Alternate version: Your impacts are too big, and they just can't be true].** As far as I can tell, "this doesn't pass the sniff test" is just a snarky way of saying, "this disagrees strongly with what I thought I knew about the world, and I am uninterested in updating that view".

For those still deciding whether or not to update their views, it seems worthwhile to explicitly lay out the key assumptions in our projection exercise, and then be explicit about why these might or might not be good assumptions.

1. *Assumptions about how much future climate might change*. Our only assumption here is that the CMIP5 ensemble is in the ballpark for RCP8.5. But that doesn't really matter that much either, since we can calculate impacts under any amount of warming you want (see Figure 5d). But people seem less worried about this one anyway.
2. *Assumptions about secular trends in growth*. Again, we pull these from the SSPs and so are just passing-the-assumption-buck, so to speak, onto what the folks who put together the SSPs assume about how countries are going to perform in the future. But these assumptions don't end up mattering too much for our main headline number (i.e. the impact on global GDP, relative to world without climate change), because that's a relative number. For comparisons between a high-baseline-growth scenario (SSP5) and a low-growth scenario (SSP3), see Extended data Table 3
3. *The assumption that historical responses will be a good guide for understanding future responses*. This is a key one for most folks, and you can see at least two reasons to be uneasy with this assumption. First, (3a) that our historical estimates are derived from year-to-year variation in temperature, which is potentially hard for agents to anticipate and respond to, whereas future climate change will be a slower-moving, more-predictable, more-anticipate-able shift. Second, (3b) that there is no way that economies 50 or 85 years from now will look like today's economies, and we can't reasonably expect them to respond similarly.

We return to (3a) below. Claim (3b) you hear a lot, and on some level has to be true: we *have no idea* what economies are going to look like in 2100 (I'm still trying to figure out what Snapchat is...). However, what we do have is the experiment over the last 50 years of being able to look at countries at very different points in the development process, and study how both the really advanced ones and the really poor ones respond to environmental change. And from our vantage point, the news is just not that good: sensitivity to temperature fluctuations has not changed over time (Fig 2c in the paper, reproduced below), and rich countries appear only marginally less sensitive -- if at all -- than poor countries (Fig 2b and like half of the supplement). This latter result, as we highlight in the paper, is very consistent with a crap-ton (technical term) of micro-level studies from rich countries -- e.g. see Sol and Tatyana's nice paper on the US. Even incredibly technologically advance countries, and advanced sectors within those countries, are hurt by higher temperatures. I just don't see how you can look at these data and be sanguine about our ability to adapt.

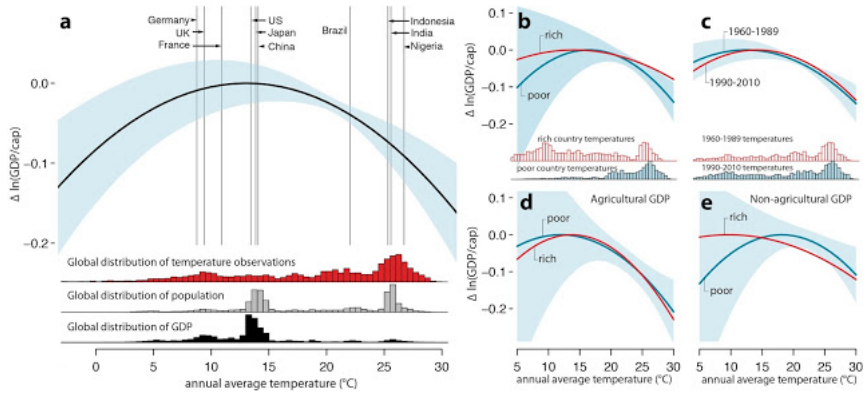

Effect of annual average temperature on economic production. **a**, Global non-linear relationship between annual average temperature and change in log gross domestic product (GDP) per capita (thick black line, relative to optimum) during 1960–2010 with 90% confidence interval (blue, clustered by country, N= 6,584). Model includes country fixed effects, flexible trends, and precipitation controls. Vertical lines indicate average temperature for selected countries. Histograms show global distribution of temperature exposure (red), population (grey), and income (black). **b**, Comparing rich (above median, red) and poor (below median, blue) countries. Blue shaded region is 90% confidence interval for poor countries. Histograms show distribution of country-year observations. **c**, Same as b but for early (1960– 1989) and late (1990–2010) subsamples (all countries). **d**, Same as b but for agricultural income. **e**, Same as b but for non-agricultural income.

So, yes, the future world might look different than the current world. But saying that is a cop-out, unless you can tell a convincing story as to exactly why the future is going to look so different than the past. Our guess is that you are going to have a hard time telling that story with an appeal to the historical record.

2. **You're studying weather, not climate.** An old saw. In fact, at every single conference related to the economics of climate change that I have ever been to, if 5 minutes passed without someone mentioning weather-versus-climate, then with probability 1 someone else would smirk-ingly mention that it had been at least 5 min since someone mentioned weather-versus-climate.

Joking aside, this is still a really important concern. The worry, already stated above, is that our historical estimates are derived from year-to-year variation in temperature ("weather"), which is potentially hard for agents to anticipate and respond to, whereas "climate" change will be a slower-moving, more-predictable, more-anticipate-able shift. Whether people respond differently to short- versus longer-run changes in temperature is an empirical question, and one that is often tricky to get a handle on in the data. Kyle Emerick and I have looked at this some in US agriculture, and we find that responses to slow-moving, multi-decadal changes in temperature don't look very different from responses to "weather" (see earlier blog here) -- hot temperatures are bad whether they show up unexpectedly in one year or whether you're exposed to them a little bit more year after year. Now whether this result in US agriculture extrapolates to aggregate country-level output in the US or anywhere else is unknown, and to us a key area for future work. But again, just claiming that responses derived from studying "weather" are a bad guide to understanding "climate" is not that satisfactory. Show us how long-run responses are going to be different.

3. **All you're picking up is spurious time trends.** This one is annoying. Please read the paper carefully, and please look at

- scientific process (1)
- sea level rise (1)
- seeds (1)
- story (1)
- supply (1)
- unsolicited recommendations (1)
- useful links; unsolicited recommendations (1)
- weather data (1)
- what does hadoop mean (1)
- wheat (1)

Blog Archive

▼ 2015 (22)
- ► November (1)
- ▼ October (3)
  - Climate change and the global economy
  - Paul Krugman on Food Economics
  - Angus Deaton and Commodity Prices
- ► September (1)
- ► August (4)
- ► June (1)
- ► May (1)
- ► April (2)
- ► March (1)
- ► February (4)
- ► January (4)

► 2014 (40)
► 2013 (38)
► 2012 (37)

Extended Data Table 1 in the paper (conveniently packaged with the main pdf, so there is no excuse!). Countries have been getting richer over time, on average, and the world has also been warming over time, on average. But since all sorts of other crap has been trending over time as well, it's clearly going to be hard to correctly identify the impact of temperature on economic output just by studying trends over time. So you have to deal with trends somehow.

So we try all sorts of combinations of year- or continent-year fixed effects, and/or linear or non-linear country time trends, to try to see how things hold up under different approaches to taking out both common shocks (the year FE) and trending stuff. If you're worried about "dynamic effects", we also control in some specifications for multiple lags of the dependent variable. It doesn't end up mattering too much. As shown in Extended Data Table 1, we still get a similar looking non-linear response no matter the model. [And, to be clear for the time series folks, our LHS is *differences* in log income, not log income]. If you still think we messed this up, then download our data and show us. The onus is on you at this point, and just making claims about the *potential* for spurious trends does a disservice to the debate.

4. **Who funded you?** The fact that I'm now getting these unsigned emails from anonymous gmail addresses I think means we did something right. I am mainly funded by Stanford University, who pays my salary. We had some project support from a $50k grant from the Stanford Institute for Innovation in Developing Economies. And I thank the Stanford Institute on Economic Policy Research for giving me a place to sit last year while this paper got written. Sol and Ted both work at Berkeley, so are paid by that fine institution, but neither received additional project support for this work from anyone.

5. [Added Oct 27]. **You do not account of the effects of development**. Or, verbatim from Richard Tol, "Although Burke and co notice that poorer countries are more vulnerable to climate change, they did not think to adjust their future projections for future development". Richard, please please read the paper before blogging this sort of stuff. This is EXACTLY what we do in Figure 5 panels b and d. "Differentiated response" means that rich and poor countries are allowed to respond differently, per Figure 2b, and that poor countries "graduate" to the rich-country response function in the future if their income rises above the historical median income. This does reduce the global projected impacts from -23% to -15% (see Fig 5b and Extended Data table 3) -- a meaningful difference, to be sure. But -15% is still about a factor of 5 larger than what's in any IAM.

Posted by Marshall Burke at 12:09 AM

3 comments:

Unknown October 26, 2015 at 8:58 AM

My criticism is the conjunction of #1 and #2:

When results don't pass the sniff test, the burden of proof is higher; it takes more evidence to convince me to update my beliefs. Because your research is based on weather, and not climate, it fails to meet that burden. You write, "Show us how the long-run responses are going to be different." Frankly, if you want to present a convincing argument, I think the burden is on you to demonstrate that they are going to be the same.

Reply

Michael Roberts October 31, 2015 at 11:07 AM

Hi Guys,

Really neat work. At this point I think weather-versus-climate issue is actually a small deal, and economic theory tells us as much (see my earlier envelope theorem post, actually inspired by Sol). And I saw some comments by John Reilly saying "it's just correlation", which seems like a feeble counterpoint given you're isolating fairly exogenous weather variations in fixed locations.

But here's the thing: you really are missing most of the economics here. Prices matter, and they're all washed out by your fixed effects. Prices guide trade and storage which can act as shock absorbers in the system, and be critical for the longer run consequences of temporary shocks. It's possible that your estimates under- or over-estimate impacts as a result. Your fixed effects likely attenuate the temporal effects. This was the problem with Deschenes and Greenstone's argument that including past weather in their farm profit regressions would account for storage. That can work, but only if year fixed effects are removed. And since weather is presumably random from year-to-year, there shouldn't be a problem with omitting the year fixed effects while including lags of weather.

Prices also matter for the transmission of shocks to other regions via trade. If U.S. farmers get hammered by climate change, they surely won't suffer, since prices will go up way more than quantities will go down. Instead, the countries we export too will pay much higher prices for less grain. And this, in turn, could change their behavior. Thus, economic impacts can be hugely dislocated in time and spaced from the physical impacts you measure. I think the right way of dealing with this is to include contemporaneous and lagged shocks of major trading partners in your regressions. If you can constrain the non-linear part to be the same for each shock--say, a degree-day measure motivated by your current estimates---and limit to X largest trading partners, maybe you will have enough degrees of freedom to identify something.

Finally, all this work will allow you to do a serious gut check: do the lagged and cross-country effects make sense and line up with what and how much they trade? Even if SEs are large, if this holds together conceptually, you would be on rock solid ground. And I gather it would give good 1st-order overall predictions for climate impacts, taking prices, substitution and adaptation into account. The only big thing it would miss is induced innovation, for which credible estimates are unlikely possible.

I'm going to see if students my graduate class want to give this a try. But you can probably crank it out yourselves very quickly.

Reply

Marshall Burke November 2, 2015 at 5:25 PM

Mike-
thanks for these really useful comments. We actually had ran the temperature-shocks-in-trading-partners regression and found that the main in-country affects remain relatively unaltered. It was too much to include in this paper but we are planning on

putting it out in follow-up work. Agreed that this seems really important.

The point on washing out price effects is good too. We had thought about this a little, and in Extended data table 1, we show a model where we drop the year FE (see column 7). Things don't change too much -- the temperature optimum is actually slightly lower, which would make our projections more negative -- and so we were happy and stopped there, but maybe there's more we can do.

Reply

Newer Post Home Older Post

Subscribe to: Post Comments (Atom)