Frequently Asked Questions for Burke, Hsiang, and Miguel (*Nature*, 2015)

1. What does your study find?
2. What does this mean for where I live?
3. How do you tell that economies perform worse at high temperatures?
4. Correlation doesn't imply causation, right?
5. Why would an economy behave differently when it is hotter or colder?
6. Do we see these patterns in any other kind of economic data?
7. Can we use these results to understand the "price" of climate change?
8. How does this compare with the models we are currently using to design climate policies, domestically and internationally?
9. If we grow slowly one year, don't we usually catch up soon after?
10. Why haven't we noticed this before?
11. Won't people adapt to warmer temperatures in the future?
12. Are you just getting these results because hot tropical countries tend to be poorer and less able to cope with the weather?
13. Isn't the main problem for climate change just agriculture?
14. If we become richer in the future, won't we just use more technologies to cope with high temperatures, like using more air conditioning?
15. People in the future will be richer than we are, so why should we care if climate change makes them a little poorer than they would have been?
16. Do your results explain why tropical countries are poorer today?
17. Does rainfall matter?
18. What about other future changes, like sea level rise and hurricanes?
19. Aren't projections of the future usually wrong? How certain can you be?

1. What does your study find?

We find that over the last fifty years, temperature has influenced the economic productivity of countries. For cold countries, warming up helps them perform better, up to a point. There is an optimal temperature at around 13C (55F) where economic performance peaks. Then warming above this temperature causes economic productivity to decline, with a rate that accelerates the hotter and hotter a country gets.

We have known for some time that the fundamental building blocks of our economies, such as workers and crops, show their highest level of performance at moderate temperatures. We demonstrate that our results for the performance of much larger and more complex national economies is consistent with what we already know about the performance of their building blocks (e.g. workers and crops).

Based on these findings, we calculate how global warming is expected to affect economies around the world. We find that climate change will reshape the global economy, causing a small number of cold countries to perform better and many temperate and hot countries to perform worse. On net, we project that the global economy will do much worse because of climate change, with global average incomes 23% lower in 2100 with climate change relative to without it. In addition, because some of the cooler richer countries are expected to benefit from warming and poorer tropical countries are hurt, global inequality is projected to get much worse due to climate change.

2. What does this mean for where I live?

We've actually built a website where you can look up custom results and projections that are specific to your country. Go to http://web.stanford.edu/~mburke/climate/ to obtain results and figures specific to your country.

3. How do you tell that economies perform worse at high temperatures?

We carefully examine economic data from 166 countries between 1960-2010, removing all changes in the data due to differences between countries (e.g. some are landlocked, others started out richer), gradual changes within each country (e.g. China underwent economic liberalization, Brazil has been restructuring its trade policies), and large shifts in global markets (e.g. the oil crises of the 1980?s). We then examine whether the economic performance of a specific country got better or worse when that country got hotter or cooler, year over year (while also accounting for the effects of changing rainfall). For countries that start out cold, warming up increases their productivity until they reach about 13C (55F). But once countries cross that optimal temperature, we find a stronger and stronger negative correlation between that country's performance and its temperature. The hotter a country is to begin with, the more economic damage is suffered from an additional 1 degree of warming.

4. Correlation doesn't imply causation, right?

Correct, correlation does not necessarily imply causation, but sometimes it does. The reason we are trained to remember "correlation doesn't imply causation" is because there are many ways that correlations have been misinterpreted in the past. For example, one might observe that Nigeria is both warmer and poorer than Norway, and conclude from this observation that Nigeria is poorer because it is warmer. This could be an flawed conclusion, since there are so many ways in which Nigeria and Norway differ that there is no way to reliably isolate the role of temperature: the two have different histories, different geographies, different cultures, different political systems, etc. They are not comparable among so many different dimensions that any correlation between temperature and economic performance will surely be confounded.

However, we are all used to thinking about simple experiments where we have a "treatment" group and a "control" group of similar things, such as plants in a greenhouse or subjects in a medical trial. The key to a good experiment is to start out with treatment and control groups that are very similar to one another initially, so that if we see that experimental manipulation of some factor caused something to change in the treatment group, relative to the un-manipulated control group. If the experiment is conducted properly, we can look at the correlation between treatment and the outcome of interest (maybe the height of the plant or the health of a medical patient) and learn what the causal effect of treatment was.

We can do something similar if we have a setting in which natural conditions have approximated these types of experimental conditions normally created in a laboratory. Our study takes advantage of one of these so-called "natural experiments": because temperature fluctuations from year to year are largely random, we can compare how individual societies perform under warmer-than-average temperatures to how they perform when temperatures are cooler, allowing us to plausibly isolate the particular role of temperature in economic performance. In the language of experiments, exposure of a particular society to hot temperatures is the "treatment", which we then compare to the "control" of that same society under cooler temperatures. Because a particular population (e.g. country) is very similar to itself across sequential observations (between two years, it still has the same history, geography, culture, political system, etc), if we see over and over that a country has lower economic production right when it also becomes hotter, then we can conclude that these environmental temperatures are likely causing these changes in economic performance.

5. Why would an economy behave differently when it is hotter or colder?

Multiple studies suggest that many of the fundamental elements of an economy, such as crops (see here or here) and workers (see here, here, here, here, and here), behave differently when the temperature changes. In each case, these studies suggest that there are "optimal" temperature where these various elements of the economy are most productive. In most cases, these optimal temperatures roughly line up to be around 20-25 degrees C (68-77F), with productivity declining sharply at higher temperatures. So maybe it shouldn't be a big surprise that the overall productivity of the entire economy rises and falls when we move in and out of this range.

In addition to productivity, other channels are starting to show up. Energy demand is very high on hot days in wealthy countries because we all turn on our AC units at the same time, and this costs a lot of money (Richard Tol's FUND model predict AC costs are the dominant economic cost of climate change). There are also emerging results about infrastructure performing poorly in the heat, people making cognitive errors in the heat, and car accidents increasing in the heat (maybe due to cognitive errors). There is also evidence of more health problems associated with very high temperatures (e.g. heart attacks in rich countries and malaria and dengue in poor countries). Finally, multiple studies observe that higher temperatures are associated with various types of human conflict around the world (e.g. see our review article here), no doubt leading to disruption of normal economic activity.

The macro-level effects we document in the current study are some combination of these various micro-level effects. Understanding which (if any) particular pathway dominates is certainly an important question, but not one which the current paper tackles.

6. Do we see these patterns in any other kind of economic data?

Yes. Some of our own work has examined the economic productivity of counties in the United States (see here) and found that it responds to temperature similar to the results obtained in this global analysis.

7. Can we use these results to understand the "price" of climate change?

Yes. We use our results to simulate global economic performance in a warmer ?business as usual? climate relative to performance in a historical climate and find that on average, unmitigated climate change reduces global average incomes by 23% by 2100 and widens global economic inequality. These results do not require that we know everything about how the future global economy evolves, since it is of course impossible for us to really know how politics, technology and other changes will alter the future global economy. These projections simply describe the net difference in economic outcomes in a world where politics, technology, etc change and a world where all of the factors change and we experience climate change. Thus, these numbers are an estimate of the additional economic cost of climate change, irrespective of a range of political, technological, or other economic changes that might occur in the coming century.

8. How does this compare with the models we are currently using to design climate policies, domestically and internationally?

Our results suggest the economic cost of climate change is much larger than previously thought. There are several models that are used to compute the economic costs of climate change and we directly compare our results the three leading models used to set US regulations (DICE, FUND, and PAGE). We find that our results suggest economic losses roughly 2.5-100 times larger than these widely used models. Differences are particularly stark at lower levels of warming. For instance, we show that substantial losses can occur even for end-of-century temperature increases of less than 2C.

9. If we grow slowly one year, don?t we usually catch up soon after?

Not necessarily. What matters most is whether households and firms work harder and sacrifice consumption to "catch up" with income they lost in a prior year, or if they try to maintain their lifestyle and patterns of consumption and instead sacrifice investments in new equipment, new technology, education, and new ventures. In the latter case, reductions in productivity today turn into lost production in later periods since there is less investment in the things that make economies highly productive. We test which case is more consistent with the data and find that reduced productivity today (due to higher temperatures) translates into reduced production in later years. In other words, we find that economies do not "catch up". This result is consistent with earlier work looking at temperature and growth (see here and here).

10. Why haven't we noticed this before?

Detecting small changes in growth rates is difficult, and figuring out what is driving them is even tougher. There is a lot of work in macroeconomics trying to understand what factors cause business cycles or persistent changes in growth rates, but a lot of this work is focused on policies or institutions that govern our economy (e.g. interest rates set by the federal reserve), with very little attention focused on environmental factors.

There is a long history of economists thinking about how geographic factors such as mountain ranges and harbors, as well as the climate, affect economic growth, but it was hard to study these relationships in a scientific way so the available evidence was generally considered flimsy. In the last two decades, there has been a tremendous amount of progress in the study of how policies, incentives, and institutions affect growth, and studying environmental factors was less popular. But only in the last ten years or so have a small community of researchers been advancing the methods needed to study how environmental factors might also influence macro-economies. Our work in this paper builds on a lot of these very recent innovations developed by this community, such as the excellent work by Dell, Jones, and Olken.

We find that changes in temperature can have persistent effects on economic output in both rich and poor countries. These effects are often small in a given year, making them hard notice in isolation, but over time they add up and can have huge effects. For instance, think of an airplane flying from point A to point B. How fast the airplane flies depends on a lot of things – the type of airplane, the amount of air traffic, the decisions of the pilot, etc – but it also depends on the headwind that the airplane faces. A slightly stronger headwind will likely be unnoticeable to passengers on the airplane, but can slow the plane down enough to making the flying time between A and B much longer. Our results suggest that introducing climate change to the global economy is like having a headwind when you're flying: you might not notice it in any given year, but over the long run it can slow you down dramatically.

11. Won't people adapt to warmer temperatures in the future?

They certainly might, although our results suggest that over the last 50 years we have not adapted much to the current climate that we are in, so we are not terribly optimistic about the next 50 years. It is of course always possible that future populations will get their act together in a way we haven't been able to do historically, but that's pure speculation and we could just as easily speculate that compounding or interacting effects of climate change will make us even worse at adaptation in the future. The bottom line is that the data do not provide us with any evidence that effective and profound adaptation should be expected.

12. Are you just getting these results because hot tropical countries tend to be poorer and less able to cope with the weather?

No. We check in multiple ways whether the fact that hot countries also tend to be poorer makes us confuse the effects of income for the effect of heat. No matter how we look at it, we find that being hot or cold is much more important for predicting how the economy will respond to warming. There is a small amount of suggestive evidence that rich countries fare a little better when they heat up (see Figure 2 in our paper), but these differences do not meaningfully affect our projected global impacts.

13. Isn't the main problem for climate change just agriculture?

In the past, many economists thought that impacts on agriculture was the primary pathway through which climate change would influence the economy, probably in part because crops are obviously sensitive to temperature to rainfall. However, in recent years it has become increasingly clear that other sectors of the economy are also vulnerable to climate change. For example, many labor-intensive sectors show sensitivity to high temperatures (see above), probably because individual people working in these sectors work less and work less productively when temperatures get too hot. There is a large literature in ergonomics studying these effects (see Figure 1b).

There are other ways in which high temperatures affect economies besides labor productivity. For example high temperature affect human health and mortality rates (see here and here), they cause us to spend more money on air conditioning (see here), and they cause increases in various types of human violence (here), which can disrupt economic productivity. We think it is likely that the reduced productivity of workers and agriculture are playing a central role in our results, although our findings represent the cumulative effect of all of these varying impacts recombining in the economy.

14. If we become richer in the future, won?t we just use more technologies to cope with high temperatures, like using more air conditioning?

It is certainty possible, although our results make us more cautious about this idea. If using advanced technologies insulated our economies against the effects of temperature, then we would expect (1) rich countries would be less sensitive than poor countries and (2) we should be getting less sensitive with time. We check for both of these effects in the data and find that neither appears to be true. Instead, we find that the effect of temperature is similar across rich and poor countries, and the effect of temperature today is almost identical to its effects in the 1960's (see Figure 2 in the paper).

There are a few reasons why this might be happening. First, for many parts of the economy, this problem is hard to "fix" with technology – e.g. it is hard to air condition agricultural fields or construction sites. Second, in cases where such a technology can be deployed, it might be expensive enough that it isn't used widely enough to make a big dent the performance of the national or global economy. Third, even in cases where such technologies (e.g. air conditioning) are widely used in the richest countries, they cost money to use and these costs take away from other investments firms might make if they didn?t have to spend money on technological solutions. Finally, even if we use some advanced technology to solve a specific temperature effect in a specific industry (e.g. air conditioning an office building), there are many other places in the economy and in people's daily lives where temperature effects are felt, constraining the overall effectiveness of the that investment. For instance, our bodies heat up when we are outside, so we don't cool down and reach thermal equilibrium for a few hours after we enter an air conditioned environment, causing us to lose productivity during this period even when there is air conditioning in an office.

15. People in the future will be richer than we are, so why should we care if climate change makes them a little poorer than they would have been?

That is certainly one view that many people have. It's not really our job as scientists to tell people how much they should care about their descendants' wellbeing. Nevertheless, our findings indicate that if climate change proceeds unmitigated, many people in the future will actually be poorer than they are today! The effects of climate change are so large for many hot counties that we expect 5-43% of countries to actually get poorer over the 21st century in absolute terms, i.e. poorer than they are today. So our results suggest that we can't just assume that everyone will keep getting richer.

16. Do your results explain why tropical countries are poorer today?

It's hard to say. To be precise, our results only compare how a country behaves relative to itself at different points in time. Comparing across countries is more difficult, since countries with different average temperatures also likely differ in many other dimension (history, culture, politics). This means we can't know for sure whether we should attribute the productivity differences we see to differences in temperature or differences in these other factors.

But what we can say is that, based on our results, if you had two countries that really were identical in all respects and you kept one at the optimal temperature (like the USA) and one at a much hotter temperature (like Nigeria) then after several years we would expect the country at the optimal temperature to be much richer than the hotter country. So the global pattern of economic wealth that we see today is consistent with what we'd expect if these effects played a major role in shaping economies historically. However, we can't be sure how much of what we see today is due to temperature and how much is from other factors.

17. Does rainfall matter?

We account for the effects of rainfall, to make sure that changing rainfall isn't actually the thing responsible for causing the response to temperature that we are studying. We do not find a systematic global pattern in how economies respond to changing rainfall. The effect of temperature seems to matter far more, which is consistent with many smaller scale studies. Perhaps we do not see large effects of rainfall because water is easier to store, so it is easier to cope with rainfall shortages most of the time, at least historically. Although with shifting rainfall patterns in the future, this might change. It will be important for future work to study this relationship.

18. What about other future changes, like sea level rise and hurricanes?

Our results really focus on the effects of temperature, and so when we simulate the future we are only focusing on the effects of temperature. Climate change is expected to change a lot of things in addition to temperature, and we are not capturing those effects. So if these other changes are harmful to the global economy on net, then our results will understate the costs of climate change.

In other research, we have studied the effects of hurricanes and sea level rise. If one wanted to compute the full impact of future climate change, one would need to combine results from the present study on the effect of temperature with results from these previous studies, as well as other impacts from other research teams.

19. Aren't projections of the future usually wrong? How certain can you be?

Our projections are highly uncertain, which is an important result in itself. Figures 4b and 5a show the range of projections that we obtain from our method, and importantly the future looks highly uncertain, especially for North America and the World as a whole (Figure 5a). Our results suggest that global average income could rise over 25% due to climate change or could fall by more than 60%, although both of these are extreme cases that are pretty unlikely to happen. We think the right way to interpret this is to realize that climate change creates tremendous uncertainty about our global economic future, and that it generates substantial likelihood that we could be much worse off relative to a world without climate change.

We note that projections for many regions of the world are not as uncertain as global mean incomes. For example, in Figure 4b there are many regions that have fairly narrow bands for their projections. These regions are those that are already quite hot, so in Figure 2a they are far to the right and simply "slide down to the right" in a fairly predictable way.

The uncertainty in colder regions arises because they are near the optimum in Figure 2a, so if we have any uncertainty in where exactly the optimum is we get very different results for the effect of warming on these countries – they could be going up or down the hump as they warm if the optimum shifts from right to left. Because these regions tend to be richer, they have a large effect on the global average incomes and generate the bulk of the uncertainty in the global mean impact (Figure 5a).