Climate Risk, Performance, and Human Capital: Evidence From 100-Year Floods

Abstract

The economic consequences of anticipating natural disasters are largely unknown. Leveraging an education setting, I combine cognitive performance data collected before, during, and after two natural disasters, spatial data at the 20 meter resolution, and a differences-in-differences identification strategy to find anticipated flood exposure reduces performance by 7% of a standard deviation. The effects are concentrated in lower elevations. I identify living in high-rises, newer areas, and buildings of greater quality as successful defensive measures. I identify long-term impacts through increased drop-out rates for the exposed. I detect defensive emigration those exposed move to higher ground. During anticipation, precipitation disproportionately dampens performance of those in low-lying areas. Using confidential access statistics for the flood warnings system, I find that flood concerns are associated with low-lying performance harms.
Keywords: climate change, natural disasters, warnings, floods, productivity, differences-in-differences, regression discontinuity, precipitation

1 Introduction

Climate change has seen the incidence and severity of natural disasters and environmental hazards increase (Field, 2014).[1] It would come as no surprise if worker productivity during a natural disaster all but ceases.[2] However, is there a productivity decrement in a disaster's *anticipation*? And if there is, what can be done about it?

This paper provides what I believe to be the first evidence that natural disaster anticipation has a detrimental impact on cognitive productivity. Identification comes through an institutional quirk of timing; students in my setting are tested during a freshet (spring thaw) which swells local river levels. While the extent to which the performance of adult students in a university exam setting coincides with the productivity of the adult workforce, it is a setting which offers a clean measure of productivity during a cognitively demanding and inflexibly scheduled task. Whether the freshet and concurrent release of accumulated winter precipitation results in fluvial flooding is only predictable a few days in advance, since water levels driven by a known quantity (depth of the snow pack in the river basin) and an unknown quantity (amount of precipitation to be received).

The setting allows me to speak in detail about the cognitive productivity of a large sample of fully incentivized adults treated by the *anticipation* of a natural disaster. I use high spatial resolution address data and repeated individual-level measures of cognitive productivity and compare them to, in a literal sense, their neighbors.[3] This paper provides an additional channel through which climate change may affect economic outcomes; the increasing incidence and severity of natural disasters may have an additional and currently unaccounted-for anticipatory effect on worker performance.

[1] There is a plethora of climate change and natural disaster literature, for an economists introduction to climate change see Hsiang and Kopp (2018), which specifically notes climatic distortions affecting the incidence of droughts and floods.

[2] Or that after a disaster, such as Hurricane Katrina, industries such as construction increase production (Groen and Polivka, 2008; Belasen and Polachek, 2008)

[3] The data is also particularly recent as flooding occurred during 2017 and 2019.

Strömberg (2007) tabulates the global humanitarian consequences of natural disasters from 1980-2004. During that period of time, there were more floods ($N = 2,102$) than any other natural disaster which affected the greatest number of people at 2.49 *billion*.[4] In terms of property damage, flooding is more costly to Canadians than fire, storms, or theft (Davies, 2020) Almost two million Canadian households are at a 'very high' risk of flooding (Henstra and Thistlethwaite, 2017). Data from the National Oceanic and Atmospheric Administration suggest that flooding-only events (unrelated to hurricanes or storms) in the United States since 1980 have caused over $151 billion in damages, unfortunately greater than even the $85 billion worth of damages attributed to wildfires during the same period. FEMA notes that floods can happen anywhere and just once inch of floodwater can cause up to $25,000 in damage. In the United Kingdom, the Environment Agency has remarked that two thirds of properties in England are at risk of flooding. Flooding is expected to increase globally; a contemporary baseline of 250 million people currently threatened by annual flood events is expected to increase by 110 million people by 2100, assuming RCP 4.5 and a stable arctic (Kulp and Strauss, 2019).

While the effects of exposure to flooding, other natural disasters, and environmental hazards (such as increasingly hot temperatures) are under heated study by the scientific community, comparatively little work examines the human and economic effects of their *anticipation*. Understanding the full extent of the economic consequences of natural disasters and environmental hazards increasingly caused by climate change is an important aspect of estimating the costs of a warming world, adjusting our current social costs of carbon, or conducting cost-benefit analysis to investments in disaster resiliency.[5]

[4] Other disasters include droughts ($N = 510$), which are deadlier at an estimated 559,000 dead and 1.59 billion affected. Mentioned elsewhere in the paper, extreme temperature events ($N = 237$) resulted in an estimated 68,000 dead and 11 million affected in the time frame.

[5] In the same vein as Moretti and Neidell (2011) whom exploit variation in Los Angeles harbor arrivals as an instrumental variable for ozone levels to find that accounting for avoidance behavior is important in understanding the full welfare effects from decreased pollution.

The outcome data that I use for cognitive productivity relates to 78,771 exams taken by 9,308 adult students over a the 4 year period from 2016 to 2019 at the University of Ottawa, a large, comprehensive, and research-intensive public university in Canada which operates from a main campus located in the center of the capital city. While the extent to which impacts on exam performance would also be seen in workplace productivity is an open question, academic scoring reflects a clean measure of mental proficiency which, at a minimum, seems likely to correlate with performance in a range of brain-intensive work tasks.[6] At least four features of the setting make it an ideal context to explore the current research question:[7]

First, water management in Ottawa is sophisticated and state-of-the-art, complete with active reservoir management. In an independent review of the Ottawa River flooding in 2019, McNeil (2019) notes that "Based on an analysis of the information available for all of the systems that experienced flooding in 2019, nothing points to human error or the negligent operation of water control structures as the cause of the flooding. The sheer amount of water (snow and rainfall) on the landscape directly contributed to the flooding. Measures taken by water managers everywhere were effective in reducing the magnitude of flooding and associated damages throughout the

[6] In the taxonomy of Harrison and List (2004), the setting is a natural experiment on a convenient sample - students. They explicitly note that the nature of a student subject pool is often criticized in experimental economics for sample selection (students are not representative of the population) and conditional on that selection, the estimated treatment effect on students is not generalizable. I believe I can address both of these issues. If the treatment effect of a warning only affects workers whose tasks are generally cognitive in nature, then the applicability of my estimates is not heavily diminished. The population of Ottawa and other major cities in the world are often characterized by a predominantly white-collar workforce, and the persistent global shift towards information and technology work would only serve to make the estimates more generalizable. Second, these students are most likely to become members of the workforce in that regard; in fact the University does have co-operative education, where for a semester students are given on the job training and experience with an employer at their place of business, earning a wage. Students treated by flooding are less likely to continue in a co-op program than control students. This could be either to poor performance at the co-op placement or withdrawal from the co-op aspect of the program, or both.

[7] A secondary benefit is the relative scarcity of disasters, natural and otherwise, to affect the city. Very few disasters occurred during the study period, limited to the Parliamentary Shooting on October 22, 2014 (motivated by Islamic Extremism and resulting in 2 dead, 3 injured) and the Dunrobin Tornado on September 21, 2018 (resulting in the destruction of 50 homes and without fatalities).

drainage basins."

Second, the flood maps in the City had been recently updated (and subsequently heavily contested) when the 2017 and 2019 flooding occurred. Flood maps aim to identify which areas are likely to experience flooding if fluvial (or other) water levels surpass certain thresholds. Unlike comparable countries such as the United States (through FEMA) or the United Kingdom, there is currently no central repository for flood maps in Canada; it is the responsibility of a municipality to develop and disseminate them. While the consultations between the City of Ottawa and local conservation authorities began in 2012, it wasn't until 2016 and 2018 that the flood maps for the various sections of the city were completed and released. As such, the most obvious technological protection against concerns of flooding is fully-exploited, and any effects I identify already account for that margin of adjustment. However, it is important to note that while flood maps may have existed, their dissemination is likely less-than-complete and more importantly, insufficiently detailed to be of use to homeowners (Henstra and Thistlethwaite, 2017).

Third, the setting provides good quality cognitive performance data on a large number of working age adults anticipating a natural disaster. The panel structure of the data provides observations on the same subject's performance multiple times, both within a warning period and between them, allowing for inference based on within-subject variation. This expels the effects of any time invariant unobserved characteristics of individuals that might influence performance.

Fourth, the nature and scheduling of the cognitive tasks faced by subjects are determined far in advance - making scheduling particularly insensitive to subsequent flood forecasts. The peak of the flooding also typically occurs after the majority of tasks are scheduled to be completed, but warnings occur during the task-period. Further, forecasts are worded in an intentionally vague manner in regards to whom is at risk - deliberately employing language such as "low-lying areas". This allows me to rule out selection effects due to displacement-in-time of activity in response to forecasts that could contaminate interfere with inference in other settings.

I find that living in a near-water low-lying area when a flood warning is issued reduces cognitive performance in a large, negative, and statistically significant way. In a differences-in-differences specification that uses data from pre- and post-flooding events combined with individual fixed effects, I find that flooding can reduce performance by approximately 7% of a standard deviation. Identification through a regression discontinuity design finds similar results. Examining non-linearity of the effect confirms that addresses well above the 100-year flood mark are largely unaffected by flooding (see Tables 2 and 3).

Secondary data allows me to identify possible steps toward climate change resilience. For example, in low-lying areas residents of taller apartment buildings are significantly less affected by flooding warnings than those in single or detached housing. A residence's vintage also matters, as residents of new buildings are significantly less sensitive to the warnings. Lastly, broadly defined building quality (as proxied by median rent) is also indicative of climate risk resilience, congruent with the main results of Beltrán et al. (2019) (see Table 6).

Following the disaster, I find that students who had been living in low-lying areas are significantly more likely to cease studies without graduation, foregoing the benefits of increased human capital. Further, those who lived in low-lying areas were also significantly more likely to change address following the disaster - to areas of both higher elevation and greater distance from major sources of fluvial flooding (see Table 7 and Table 8).

Within the approximately two-week exam seasons, I use treatments that vary at the daily level to identify effects using students as their own control groups. For example, greater levels of precipitation during the flooded exam seasons were associated with dampened performance. A one standard deviation increase in daily rainfall reduced performance by 4 to 5% of a standard deviation during these times (see Table 10).

I also use confidential data on the daily access statistics for the web-based flood warning system operated by the Ottawa River Regulator. Using within-student variation and the independence of the exam scheduling, I note that when low-lying address students write exams on days with higher degrees of concern, their performance is markedly lower than their higher-ground peers (see Table 14).

It is important to note that this paper does not advocate against prior warning systems. It is entirely plausible that a performance decrement is a *rational* response to flood risk, in that a student may evaluate the utility of applying their effort toward protecting their domicile to be higher than the utility from the marginal performance increase of effort toward study. While this may be the case, it highlights the presence of a negative environmental impact on a student's performance nor does it alleviate the additional distributional effects from students' differential exposures to the environmental dis-amenity. In fact, more sophisticated warning systems may be better able to determine areas that are 'at risk', thereby reducing the likelihood of costly errors (of either type):the student who is at risk but is not warned of potential flooding and the student who is warned but is not at risk. For example, in Table 15, I estimate the direct effects of a warning and find that during the first flooding most students are affected, but for the subsequent crisis only those in very low and previously exposed areas were significantly affected by the warning.

The remainder of this paper is structured as follows: Section 2 provides a primer of related literature. Section 3 provides a brief summary of flooding in Ottawa. Section 4 details the sources of data used. Section 5 details the differences-in-differences identification strategy. Section 6 presents the results of the paper. Section 7 concludes.

2 Literature

The environment is increasingly recognized as an important factor in economic productivity. This section provides a primer of some related literature to the main research question: does the threat of environmental disaster reduce productivity? I first review some relevant literature concerning how exposure to environmental hazards (e.g. hot temperatures, air pollution) affect labor productivity. I then review literature concerning the response to warnings of environmental hazards.

2.1 Environmental Hazards and Labor Productivity

The research conducted in this paper is closely related to the literature that examines the negative impacts of environmental hazards on the productivity of labor.

Hot temperatures, a primary consequence of climate change, have been connected to both reduced labor and capital productivity. Hsiang (2010) shows that economic output has a structurally similar response to rising temperatures as labor productivity. In particular, non-agricultural output losses of -2.4% per one degree Celsius were larger than agricultural losses of 0.1%. In the sample of 28 Caribbean countries, he also examines the effects of tropical cyclones - an important correlate with temperature. Losses are approximately 1% for every 1 standard deviation increase in cyclone exposure - except for the construction sector which sees a 1.4% increase. Dell et al. (2012) examine country level GDP growth rates from 1950-2005. They find that poor countries suffer a reduction in economic growth of -1.39 percentage points for every one degree Celsius net rise in annual temperature, hypothesizing this to be either the result of a temperature-sensitive agricultural sector or the response of human labor to thermal stress. Graff Zivin and Neidell (2014) exploit plausibly exogenous variation in daily temperatures within United States counties to find that increased temperatures cause large reductions in labor supply (as measured in the 2003-2006 American Time of

Use Survey) for workers in industries with high exposure to climate (as defined by the National Institute of Occupational Safety "heat-exposed industries" 1986). They find monotonic declines in labor supply above 85 degrees Fahrenheit (29.4 Celsius). Zhang et al. (2018) combine weather data with 1998-2007 production data from half a million Chinese manufacturing plants. They find that both labor intensive and capital intensive firms (light industry such as food processing and heavy industry such as non-metallic minerals, respectively) are both sensitive to high temperatures - an extra day with temperatures above 90 degrees Fahrenheit (32.2 Celsius) decreases output by 0.45% relative to a baseline day of 50-60 degrees Fahrenheit (10 to 15.6 Celsius).

More specifically, hot temperatures are well studied to be detrimental to cognitive productivity in a literature that commonly leverages exam performance data for the same reasons as the present research. Zivin et al. (2020) combine weather data with test scores from the 2005-2011 National College Entrance Exams in China. They find that a 3.6 degree Fahrenheit (2 Celsius) increase in temperature during the exam period decreases total test scores by 5.83% of a standard deviation, with effects concentrated on the *highest* performing students. Park (2020) combines local daily weather data with 1998-2011 test scores of New York City secondary school students taking high stakes exams over the course of several days each year. Exploiting quasi-random variation in temperature (using a student fixed-effect model), he finds that hot temperatures on the day of the exam can reduce performance by up to 13% of a standard deviation (specifically, standardized Regents exam score falls by 0.009 to 0.011 standard deviations for every increase in degree Fahrenheit). and have more persistent effects such as impacting high school graduation status. Cook and Heyes (2020) conduct a similar exercise, combining local weather data with 2007-2016 university exam scores. They find particularly cold temperatures also have a negative effect on cognitive labor productivity.

Another environmental hazard, pollution (in particular air pollution) has also been connected to reduced cognitive labor productivity. Archsmith et al. (2018) use within-worker (United States Major League baseball umpires, which travel excessively for work that is scheduled well in advance) variation in exposure to air pollution and find that a ten unit increase of $PM_{2.5}$ increases the propensity of workers to make visible and uniquely verifiable errors by 2.6% in the 2008-2015 seasons. Returning to using exam performance data, Ebenstein et al. (2016) exploit the same student taking multiple exams, and pair transitory particulate matter ($PM_{2.5}$) exposure to scores on 2000-2002 Israeli post-secondary matriculation exams. They find a robust negative relationship between pollution exposure and test scores; a ten unit increase of $PM_{2.5}$ (44% of one standard deviation) is associated with a 0.55 points decrease (2.3% of a standard deviation) in a student's test score.

2.2 Possible Channels Through Which Disasters Affect Labor Productivity

To the best of my knowledge, the effects of exposure (or risk of exposure) from natural disasters such as flooding on labor productivity are relatively unstudied in economics. However, other academic fields have examined possible channels through which labor productivity could be affected, in particular increased aggression and anxiety.

Scott et al. (2014) conduct a psychological study to analyze 191 fourth to eighth graders who experienced Hurricane Katrina. Through estimated increases in 'aggression', academic achievement was found to be negatively impacted by previous disaster exposure. Meier et al. (2010) from studies the effects of Hurricane Katrina and Hurricane Rita on fourth to eighth graders, and estimates that Louisiana students that were *displaced* into Texas school boards are less likely to perform well on high stakes exams, although they must use pooled Texas school board data to identify these effects. The primary difference between the above studies and the current one is that, in the current study, cognitive performance of subjects is measured during the disaster's warning period (identifying the warning, rather than exposure effect) and after (identifying the warning *plus* exposure effects). Another primal difference is the construction of a *local* control group who, for all intents and purposes, would be identical to the treatment group absent treatment.

This paper contributes directly to the flood risk and mitigation cost-benefit literature, chiefly advanced by hydrology engineers. In particular, I examine the effects of flooding on labor productivity and human capital accumulation, 'intangible' metrics (so called by engineers as they are difficult to put into monetary terms) which have proven difficult to measure for decades Lekuthai and Vongvisessomjai (2001); Dassanayake et al. (2015) either due to a lack of data or a lack of appropriate evaluation methodologies. Lekuthai and Vongvisessomjai (2001) represent the first attempt, wherein the authors examine how flood severity affects survey-reported anxiety levels in Bangkok. They find that a small proportion of the population felt that they would be affected by a 'low' level flood and a greater proportion was concerned when a 'high' level flood was hypothesized. In this paper, I advance the rigor of this literature toward revealed, rather than stated, effects.

2.3 Responses to Environmental Hazard Warnings

This paper is also related to a literature which examines the effects of *warnings* of environmental hazards on, admittedly varied, economic outcomes.

Saberian et al. (2017) study the effect of air pollution warning systems on avoidance behavior, namely reductions in Australian outdoor cycling during 2008-2013. In their setting, air quality alerts are established the day prior to being issued and are based on the previously *forecast* air quality. They find that, when an air quality warning is issued, the amount of cycling observed is reduced by 14-35%. Rabassa et al. (2020) find similar results wherein female and older users are less likely to use a bike sharing service following a heat warning. Zivin and Neidell (2009) use turnstile data from

1989-1997 to show that air quality alerts impact attendance at two popular outdoor venues in Southern California (Los Angeles Zoo and Botanical Gardens and Griffith Park Observatory). They find that attendance declines by 15% and 5% on the day of the alert, respectively. In a model capturing two consecutive days, they find no effect of avoidance behavior, suggesting consumers disregard warnings if on consecutive days. Keiser et al. (2018), in contrast, document that visitation to United States national parks are relatively insensitive to prevailing ozone conditions from 1990-2014. Beatty et al. (2019) examines how government information warning households to acquire emergency supplies before hurricanes make landfall between 2002 and 2012. They combine forecasts with sales (derived from checkout scanner data) on bottled water, batteries, and flashlights. The bulk of sales occur immediately prior to predicted landfall - with preparation higher in coastal, wealthier, and whiter areas. Similarly, Yu et al. (2019) examines the effects of Taiwanese typhoon information on the daily vegetable wholesale market. During the warning period immediately prior to a typhoon, the market price significantly increases with no associated change in quantity traded. When the typhoon makes landfall the market price increases again with a fall in quantity traded. The effects of storm warnings persist for storms up to five days after the warning is lifted. Barwick et al. (2019) study the release of 2013-2014 real time air quality monitoring and disclosure program information in China. Exploiting staggered implementation to identify the effects of the program, they find air quality information disclosures caused households to increase searches for pollution-related topics, behavioral changes to reduce exposure, and a higher willingness to pay for housing in less polluted areas. The authors estimate that simply due to increased information access, the mortality cost of air pollution fell nearly 7%.

This paper contributes to this robust literature by explicitly examining how cognitive productivity is affected immediately prior to a forecast event - that is - I can observe young adults performing cognitively strenuous tasks during a period of time that they are warned of an impending environmental hazard, and compare their performance to those to whom the warning is less applicable.

A paper which is close to the current research is Beltrán et al. (2019), whom use a repeat-sales model to analyze house sales prices of properties affected by flooding in England between 1995 and 2014. They find that immediately *after* a flood, an affected property postcode sells for less than a non-flooded property. This discount is short-lived with the exception of lower-priced properties. Flats (apartments) have a smaller discount than semi and detached housing. For in-land (as opposed to coastal) flooding, areas with defenses experienced less of a discount. To address identification concerns - flooding most often occurs to houses on the *flood plain* - they restrict their analysis to properties that are on the plain itself, bearing the same risk but were not actually exposed to flooding. The discount for properties that are flooded repeatedly is not different than the discount on properties that were flooded only once.

The closest paper to the current research is Hossain (2019), who constructs a panel data-set of flood inundations to examine their impact on the Indian manufacturing sector. She estimates that a 10% increase in the incidence of flooding causes a 17.3% reduction in aggregate productivity. She also documents heterogeneity in vulnerability to flooding - the lowest productivity firms are the greatest affected. Flooding is identified using a well-proven algorithm using satellite images to overcome data availability challenges, allowing her analysis to be conducted at the 500m $\times$ 500m resolution. While satellite imagery presents a solution to the dearth of flooding data in the developing world; my setting is data rich in that I am able to leverage a 20 meter $\times$ 20 meter resolution elevation map to determine the likelihood of flooding. This level of spatial resolution is critical when comparing the productivity of neighbors.

A framework which is useful for thinking about natural disasters and their risk is well presented in Strömberg (2007). In a natural disaster there is first a triggering 'natural hazard event', which in this setting is fluvial flooding. Second is the population exposed to the event, in this setting more than 15% of addresses are below the flood level mark and close to the rivers, as depicted in Figure 5. Third, the vulnerability of the population to the hazard. In this setting, flood prevention is considered state-of-the-art, however up to 60% of the river basin is unregulated (without significant reservoirs that can be used to absorb rapid increases in water volume) leaving water levels vulnerable to amount of precipitation.

3 The 2017 and 2019 Floods

Ottawa has two bodies of water that have the potential for flooding. The Rideau River has a basin covering 4,000km^2 and discharges approximately 35 cubic meters per second into the Ottawa River, which has a basin of 146,000km^2 (a land area twice the size of New Brunswick) and can discharge more than 7,000 cubic meters per second during a normal freshet.[8]

The Ontario and Quebec governments (the river serves partially as their provincial border) created a regulatory board responsible for issuing flood warnings to the general public and government agencies. Commissioned in 1983 they are the Ottawa River Regulating Planning Board and the Ottawa River Regulating Committee, respectively. According to these bodies, flooding in the area can only be predicted a few days in advance (Ottawa River Regulation Planning Board, 2017), and a system of timely flood warnings represents the state of the art for temporary flood protection and mitigation of damage, injury, and loss of life (Parker, 2017).

[8]A river basin is any area of land where precipitation collects and drains into a river. They can be partitioned into smaller sub-basins, often with unique hydrological properties. During the Canadian winter, most precipitation is simply stored as snow or ice on the ground. During the spring thaw, large quantities of water are released and combine with rainfall, which can lead to heavy spring flow and flooding. This is called freshet.

3.1 2017 Flooding

In this section, I provide a summary of the causes and consequences of the flooding that occurred in 2017. Additional details can be found in the summary published by Environment and Climate Change Canada, accessible here, and the Canadian Broadcasting Corporation has published a photo-timeline, which I have archived and made available to the reader here. The 100-year flooding in Ottawa in 2017 was 'ranked' the year's third most significant Canadian weather event by Environment Canada, following only the British Columbia wildfires and a drought affecting the Western Prairies.

During the freshet period (April and May) 174% of the normal amount of precipitation fell in the Ottawa River Basin. Approximately 159 mm of rain fell in April alone - the largest amount in 125 years. This was followed by an additional 140 mm of rain in the seven day period from April 30 to May 6. It is often assumed that the snow pack depth in the Ottawa River basin would be enough to predict the incidence and severity of flooding. Unfortunately, this is not the case. For example, the snow accumulation was actually higher in 2016 (when there was no flooding) as compared to 2017. Snow pack is best described as a risk factor for flooding; other considerations such as precipitation and temperature (which determines the rate of thaw) are more prominent drivers of flooding.

The precipitation in the basin drains directly into the Ottawa River, with greater volume of runoff causing greater river swelling. Note that much of the river basin near the city is unregulated, that is without significant reservoirs to retain runoff (and accounts for approximately 60% of the watershed). In areas that are regulated, reservoirs were empty at the end of March in preparation for the freshet. During the flooding, the reservoirs were being used almost to capacity, and staff could not drain them to normal levels until much later than usual. It is estimated that without the installation and operation of the 13 available reservoirs for Ottawa River management, peak water levels could have been up to 90 cm higher. Since the peak above-datum measurement was 2.54 meters, this represents an avoided increase of approximately 35% of peak water levels, and consequently, avoiding more than proportionate damage.

The river level peaked at 60.44 meters on May 6, 2017, (measured in the west end of the city at the 'Britannia' monitor), which set a one hundred year historic record for levels unseen since 1915.

During the flooding, the regulator issued three press releases - April 5, 18, and 28. This coincides with before, midway and at the end of the exam period. Below are excerpts.[9] I have also provided a timeline of water levels, the dates of the below statements and the exam period, in Figure 3.

> `April 5, 2017:` The Ottawa River Regulating Committee cautions that level and flows on the Ottawa River are expected to significantly increase this week due to warming temperatures and rainfall. ... Officials of the Committee would like to inform residents in low-lying areas along the shores of the Ottawa River... of the possibility of minor flooding within the next few days depending on the amount of precipitation received.

> `April 18, 2017:` ... recent rainfall combined with snowmelt has caused the levels and flows on the Ottawa River to rise rapidly in several areas ... Officials of the Committee would like to remind residents and communities located in areas that are flood-prone to continue to monitor river conditions... levels possibly reaching 59.97 m at Britannia... Further bulletins will be issued if the situation changes.

> `April 28, 2017:` ... recent peak levels not seen in the last 20 years in many locations... additional precipitation is expected to once again increase levels that had been in decline ... Residents who would like to report flooding or require help should communicate with responsible provincial agencies.

[9] While the press releases are no longer hosted at the Ottawa River Regulation Planning Board website, I have made them available here.

The following year, 2018, there were no serious floods nor flood warnings issued. It wasn't until the end of April that the following *laissez-faire* bulletin was published by the regulator:

> `April 29, 2018:` ... based on the current snowpack and the weather forecast, levels and flows are expected to remain within the normal range of fluctuations associated with this period of the year.

According to the Insurance Bureau of Canada, the 2017 flooding resulted in 15,750 claims and $223 million in property damages. More than 5,000 residences were flooded and 550 roads were washed or swept away.

3.2 2019 Flooding

The 2019 Flooding in Ottawa was far worse than the 100-year flood experienced two years prior. This time, Environment Canada ranked it as the most significant weather event of the year. Meteorological conditions created a 'perfect storm' for excessive fluvial flooding. First, a drawn out winter and below-normal temperatures kept the ground frozen and thawed late in the season, ultimately reducing the ability of the ground to absorb its customary proportion of typical runoff. This little thawing resulted in a snowpack which was 50% greater in volume than normal. Second, twice the normal amount of rain (150mm) fell between April 10 and May 10, exactly as seen in 2017 but now with the additional volume of water from thawing.

On May 1, the Ottawa River rose 30 cm beyond the peak seen in 2017. Dozens of other rivers in the region also broke records. More than 6,000 dwellings were flooded or at very high risk of flooding. It was not until June that the City of Ottawa was able to lift its state of emergency. During that time the Regulator issued numerous statements, excerpts below. I have also provided a timeline of water levels, their dates, the below statements, and the exam period in Figure 4. The dates of the press releases coincide with the beginning, middle, and penultimate day of the exam season.

> `April 11, 2019:` ... with the onset of the spring freshet period ... Minor flooding in low-lying areas may occur over the next few days depending on the temperatures and precipitation amounts received. ... **will issue further bulletins if flood risks increase.**[10]

> `April 16, 2019:` ... With significant rainfall forecasted later this week, levels will exceed minor flood levels ... in flood prone areas along the Ottawa River ... **They will however remain below the maximum levels that were reached in May 2017**.[11]

> `April 18, 2019:` ... flows and levels similar to those observed at the height of the May 2017 flood may be reached depending on the amount of precipitation received. All flood-prone areas along the Ottawa River ... are at risk.

> `April 25, 2019:` **On Monday or Tuesday, flows and levels are expected to exceed those observed at the height of the May 2017 flood, depending on the amount of precipitation received.** All areas along the Ottawa River that were impacted in 2017, and possibly additional areas, are at risk.[12]

Two more statements were released in May 2019, after the examination period finished. They warn of even worse flooding and catalog the never before seen water levels throughout the region. It is estimated that the optimized use of the reservoirs reduced peak water levels by at least 40cm. As the peak was 2.8 meters over datum, this represents an avoidance of 14%.

[10] Boldface in source document.
[11] Boldface mine.
[12] Boldface in source document.

4 Data

In this section, I detail the data sources used in the paper. First, I describe the primary outcome data - the performance of fully incentivized adults conducting cognitively strenuous tasks that are scheduled well in advance - university exams. Second, I describe the algorithm and application resources required to accurately geocode privacy-protecting postal codes into high resolution spatial data. This includes not only latitude and longitude, but leverages a 20 meter resolution elevation mapping provided by Natural Resources Canada. A number of other data sources, which allow for supplemental analyses, are also described herein.

4.1 Performance Data

I use administrative data from the University of Ottawa as a measure of cognitive performance. In particular, I observe the universe of undergraduate student transcripts for students beginning studies anytime after September 2007 and until April 2019. To limit sample bias, I elect to restrict the sample to a time frame meaningfully similar to the years with flooding. For the 2015-2019 period, 710,492 grades were achieved by 36,518 distinct students. I connect this performance data with institutionally provided student information such as gender, age, and address via postal code. The postal code was chosen as a compromise between student privacy and accuracy in determining treatment status.

Nevertheless, the six character Canadian postal code is a surprisingly high resolution proxy for address. In an urban setting like Ottawa, a postal code can be unique to a single apartment building. In Figure 2, I present the postal code density in an Ottawa neighborhood. The average distance between a postal code and its nearest neighbor is 42 meters. I use the student postal code to construct my estimation sample of students who are plausibly at risk of flood exposure and a reasonable control group, which totals

9,308 students completing 78,771 courses.

The academic year is split into two semesters. Fall semester courses are taken from September through November, with final exams written in December. Winter semester courses are taken from January through March, with final exams written in early April. Because of my interest in the spring freshet and associated fluvial flooding, I focus primarily on exams written in April. To fix time periods, exams began on April 11 in 2017 and on April 9 in 2019, see Figure 3 and Figure 4 for a graphic representation.

Although focusing on adults taking undergraduate courses limits the external validity of the analysis, my administrative data provides several advantages over reasonable alternatives used in related literature such as surveys (such as in Graff Zivin and Neidell (2014)) or other measures of workplace productivity (such as in Archsmith et al. (2018)).[13] First, students have a very limited capacity to reschedule exams, and even more-so due to flooding concerns (particularly as it had yet to occur and was, at the time, also uncertain), allowing for the identification of productivity decrements without the potentially confounding effects of rescheduling. Second, student data allows for a clean linkage between performance at place-of-work and their address, which is arguably absent in many a restricted public use microdata file. Because this paper's focus is on flooding risk *at home* affecting performance *at work* this is a key element to the analysis. Third, student performance data constitute a panel which is available over a long period of time (providing within-student variation) and space (between student variation). Lastly, the performance testing is fortuitously timed with the spring freshet, and so I examine these students' productivity in their natural work environment during their customary timing.

Using course-level grades as the measure of performance introduces a complication. While the hypothesis is that anticipated flood risk exposure impacts performance during the final exam period, the assessment of a course is based only partially on

[13]Not the least of which is that administrative data has the benefit of being free from recall errors that arise in survey data as noted in Zivin and Neidell (2009).

the final exam. Often, other elements such as midterms and coursework completed during the semester also contribute. Academic regulations at the university require that a final exam comprise no lower than 40% and no higher than 60% of a course grade. Any variation in weighting here would only add measurement error to the dependent variable and is unlikely to be correlated with my regressor of interest - flood risk exposure. An additional source of measurement error is the granularity of course grade reporting; they are recorded as letters which correspond to a score interval. For example, an 'A' corresponds to a score in the interval 85-89%, which I summarily assign the midpoint of its interval. While measurement error in the dependent variable does not bias ordinary least squares estimates, it does increase the estimated standard errors. This has the byproduct of making any of the analysis' claims of statistical significance conservative (rather than optimistic). Additionally, I apply a multiplier to reflect that any impact (environmental or otherwise) on final exam performance has a dampened impact on course-level performance through the mechanical relationship that course work completed during the semester also contributed to the course grade. Throughout, I multiply any exam-day estimated effect (such as flood risk) on exam performance as a factor of two times the effect in course performance, consistent with the assumption of a final exam carrying 50% of the weight in every course. In doing so, an estimate of a 3.5% decrement in overall course performance will be reported to be a 7% decrement in final exam performance.

4.2 Census Data

I connect the performance data six character postal code to the 2016 Census (Product Code 98-401-X2016044) using the Postal Code Conversion File (from the Canada Post Corporation) which provides a mapping between the six character postal code and Statistics Canada's Dissemination Areas. Postal codes are typically smaller than dissemination areas, and this is reflected in the data. For the 2016-2019 period, the

36,518 distinct students represent 14,872 unique postal codes nested within 3,494 dissemination areas. These dissemination areas, although offering less spatial resolution, remain quite small. For areas within 20 KM of campus, there is an average of 672 individuals inside an area 0.78 square kilometers.

4.3 Geolocation, Elevation, Distance to Water

The first step in identifying potential flood exposure is to *geocode* postal codes, which is the process of converting addresses (like that of the main campus `75 Laurier Avenue, Ottawa, ON K1N 6N5`) into geographic coordinates as latitude and longitude (45.423144,-75.683138). In order to do this, I use the application programming interface of the OpenCage Geocoder to identify the 14,872 postal codes provided by the performance data. The resulting coordinates are consistent with those using navigation services such as Google Maps.

Another method of geocoding the data could be to use the Postal Code Conversion File which provides an associated latitude and longitude for a six character postal code. Unfortunately, Statistics Canada has found that the median location error using the PCCF is approximately 120 meters (compared to that of geocoded results) which is far too large in the current setting which relies explicitly on spatial accuracy and resolution.[14]

Once each address is accurately geocoded, I feed the latitude and longitude coordinates into Natural Resource Canada's Elevation API. This service requests information from the Canadian Digital Surface Model, an approximately 20 meter resolution map of Canada. Combined with the latitude and longitude from the geocoder, I now have a dataset in three dimensions, which can approximate potential flood exposure using

[14] Statistics Canada analyzes the measurement error using postal codes versus full street addresses. Their Table 2 notes that the median distance between full street address and postal code conversion file is 160 meters for "ordinary urban households" which account for more than 85% of addresses in Ottawa.

distance to water, and the height relative to it.[15]

Distance to water is calculated as the shortest (non-euclidean) line that can be drawn from an addresses' coordinates to the nearest edge of the nearest body of water. To do this, I use a GIS shapefile used by City of Ottawa hydrological services which outlines any major or minor bodies of water in the region. These include both the North-flowing tributary Rideau River and the significantly larger Ottawa River (which flows West to East near the city).

4.4 Water Levels

The Ottawa River's water level is measured in meters above sea level. Environment Canada provides real-time information on the water level at various points. For the current research paper, the closest monitoring station is located in the West end of the city and is so to take advantage of the fact that the river runs West to East in the area. Specifically, I use data from monitoring station ID number 02KF005 which is the station reported by the Regulator (and the community) when referring to water levels near the city. For this station, the threshold above which the water level is considered a minor flood (a once in a 2-year flood) is 59.42 meters. The next threshold, which establishes a major flood (20-year) is 59.98 meters. The exceptional flood threshold, which corresponds to a 100-year flood, is 60.24 meters.

The 2017 and 2019 floods were exceptional, or once-in-one-hundred-year floods. I depict just how exceptional the peak levels were in Figure 1. In the thirty years prior to the most recent flooding, annual water levels never crossed the 'major flooding' threshold.

In Figure 3, I present the water levels observed at the monitoring station during the 2017 spring freshet. The blue connected dot series represents the water level. Open

[15] Using elevation data to estimate flood exposure is often the state of the art in this space - see for example Kulp and Strauss (2019). It is also possible to create satellite imagery derived flood mapping, such as in Hossain (2019), but the additional assumptions and lower resolution make the approach less attractive in this space.

data markers on April 11, 12 and May 6, represent changes in the flood status. For example, on April 11 (the first day of exams) the rising water level triggered the first stage of flooding readiness described as 'monitoring'. It was only the following day that minor flood levels were observed. Water levels seemed to be decreasing until just after the April 28 regulator announcement, when precipitation was predicted to reverse the decreasing water levels. This was proven correct; shortly after on May 6 (well after exams were completed) major flood levels were observed.

In Figure 4, I present the water levels observed at the monitoring station during the 2019 spring freshet. The blue connected dot series represents the water level. Open data markers on April 19, 20, 24 and 26, represent changes in the flood status. On April 19, the eleventh day of exams, the rising water level triggered the first stage of flooding readiness described as 'monitoring'. Once again, the following day minor flood levels were observed. On the sixteenth day water levels had risen to major flood. On the last day of exams the floods were classifiable as an exceptional or 100-year flood, the second one since two years earlier in 2017.

4.5 Disaster Recovery Assistance Ontario

There is no publicly available or comprehensive list of addresses affected by either flooding. However, following a natural disaster, the Ontario Government can implement a 'Disaster Recovery Assistance for Ontarians' program which was activated for Ottawa River flooding on May 5 2017 and May 7 2019. I identify which postal codes were eligible for disaster relief (and by proxy those that are more likely to have been affected) by manually transposing the publicly available maps onto geocoded Canadian postal codes to determine which addresses were eligible for assistance and which ones were not. An example of this manual mapping can be seen in Figure 2. A total of 437 postal codes were identified in this way for 2017 whereas 3,766 postal codes were identified in this way for 2019.

All of the addresses identified in the 2017 flooding recovery were also identified in the 2019 recovery area. These addresses are limited to 97 meters in elevation and 1,812 meters from the nearest water source.

4.6 Weather

Weather data was collected from Environment and Climate Change Canada *in situ* monitors. Information was collected from two sites. The Ottawa International Airport (ID number 6106001) provided hourly data on temperature and weather conditions (such as partly cloudy, drizzle, snow, and rain). The Experimental Farm (ID number 6105976) provided detailed daily precipitation records and snow-on-ground measures.

4.7 Summary Statistics

Summary statistics relating to course performance, student characteristics and geolocation are presented in Table 1. I use a total of 78,771 exams, written by 9,308 students.

The average course grade is 75.603% corresponding to a 'B+' in the alphanumeric grading scheme. Grades vary considerably both within and between student, while the overall standard deviation is 14.30, the within student variation is 10.29, or two letter grades around the mean.

The majority of classes are taken by full time students at 89%. Students are young, with a mean of 21 years but a standard deviation containing 17 years old (age exhibits a distinct right-skew). Female students take more than half of courses and foreign students (defined either by immigration status or by paying 'international' student fees) take just over one quarter of courses.

Mean elevation is 70 meters - with a one standard deviation reduction below the flooding water mark of 60 meters. A histogram of address elevation is also provided in Figure 5. A vertical red line is provided for elevations unlikely to have been directly affected by flooding. Distance to water is small, with an average of 474 meters to the nearest body of water. A histogram of distance to water is presented in Figure 6. Distance to the river is relatively uniform, however, artifacts of city planning can be seen in spikes above 200, 400, and 600 meters away from water. Students, unsurprisingly, live close to the campus at an average of 5.4 kilometers away. The average of the median rent of dissemination areas a student resides is $980.

5 Identification

In this section, I detail the identification strategy used to estimate the causal impacts of flood risk warnings on cognitive performance. For analyses that use pre-period data (e.g. 2016 freshet exams) and treatment period data (2017 freshet exams), identification comes from the independence of flood year status (e.g. 0 if not flooded or 1 if flooded) from student enrolment and the scheduling of exams. This assumption seems reasonable in the present context - if it were possible to predict future natural disasters with the certainty required to invalidate the assumption, the consequent damages of the disasters would be far smaller than observed.

Winter semester exams are held in an exam period that runs from early to late April coinciding with the beginnings of a spring freshet.[16] The earliest day-of-the-year exams began was April 9 while the latest exam period began on April 15. Exams are held in one of three time slots (beginning at 9:30 am, 2 pm and 7 pm). [17] The exam schedule is determined approximately six weeks before the exam period begins, well

[16]For the winter 2016 semester classes began on January 11 and ended on April 12. Exams began on April 14 and ended on April 27. For the winter 2017 semester classes began on January 9 and ended on April 8. Exams began on April 11 and ended on April 28. For the winter 2018 semester classes began on January 8 and ended on April 11. Exams began on April 13 and ended on April 26. For the winter 2019 semester classes began on January 7 and ended on April 5. Exams began on April 9 and ended on April 26.

[17]I do not observe students allowed to defer an exam to a date other than that mandated for the course, typically about 4% of the total. Deferment for reasons unrelated to future flooding (family bereavement, religious holiday, etc.) are not a threat to identification. Insofar as some deferments result from future flooding it is plausible that it works against the direction of any effect that we find, since postponement from writing an exam during a flood is likely to be to a less stressful date. However this is a valid caveat to hold in mind. The university as an administration never canceled exams during the study period.

before flooding can be forecast. Further, the last date to un-register for courses without academic penalty is mid-march, well before temperatures have warmed in earnest and the speed of thaw (a risk factor for fluvial flooding) can be determined.

5.1 Differences-in-Differences

In this section, I detail a spatial differences-in-differences identification strategy estimated by ordinary least squares. The classic DID has a treated and control group observed in a pre-treatment and post-treatment period. For this paper, the treated group is defined as having an address with an elevation below 60 meters (just below the absolute maximum peaks achieved), while the control group resides above 60 meters. This delineation is meant to reflect the flood communications by authorities that those in 'low lying' and 'flood prone areas' should monitor water levels closely.

Both the treatment and control group are required to be within 1km of a significant water body. I further enrich the model with individual student fixed effects. The fact that flooding occurs *twice* in my study period could present a complication when assigning pre- and post-treatment status, particularly in a differences-in-differences framework.[18] To conduct a more straightforward analysis, I compare, within-student, whether their pre-treatment performance in winter 2016 exams is significantly different than their post-treatment performance in winter 2017 exams, by treatment group. My specification is:[19]

$$Grade_{i,f,t} = \alpha + \gamma * Treated_f + \lambda * Post_t + \delta * (Treated_f \times Post_t) + \nu_i + \varepsilon_{i,f,t} \quad (1)$$

[18] If I considered 2016 and 2018 as pre-treatment and 2017 and 2019 as post-treatment in the same regression, then students who migrate between the two groups after the 2017 floods could act as individuals in the control group for 2017 - see the emerging literature on DID with varying time-treatments.

[19] I have deliberately chosen to follow the coefficient naming conventions in Angrist and Pischke (2008), see page 233.

Where $Grade_{i,f,t}$ is the imputed exam performance for individual i in area f completing an exam in period t. My parameter of interest is δ, the coefficient coinciding with the performance change of those warned of likely flooding once the flood warning has been issued, beyond the magnitude reported by γ, which reports the average difference of the treated and control group in the pre-period. λ identifies the average change in performance of students in the control group from the pre- to treatment-period. The inclusion of student (ν_i) fixed effects implies that identification comes from within-student variation. In other words, variations in the performance of individual subjects under alternative flooding treatments. Said differently, the interpretation of δ is the average effect on grades of a treated group student moving from the pre-control condition to the treated condition (being warned of flooding). I explore the robustness of my estimates using alternative measures later. The standard errors are clustered at the student level. Later, I demonstrate that my results are robust to a number of other plausible clustering strategies.

6 Results

In this section, I present the main results of the paper. First, I identify a productivity decrement from possible exposure to flooding. I then examine some of the possible mechanisms at play, or, alternatively the defensive investments available to reduce the harm of potential flooding in the future. I then examine the non-linearity of the effect, and find the results are concentrated, unsurprisingly, more in lower-elevation addresses. Second, I present evidence to a longer-run economic consequence; students in low lying areas are subsequently much more likely to drop out of studies altogether. Third, I present the post-flooding results of flood plain emigration following flooding. Exposed students move, on average, to an elevation 2.3 meters higher than their flood-period elevation which coincides almost directly with the additional water level seen at both flood peaks.

6.1 Performance

In Table 2, I analyze how cognitive performance (proxied by achieved exam grades) of students living in low-lying and flood prone areas are affected by the risks of flooding.

The dependent variable is standard deviations of final exam score. The main independent variable $Risk = 1 \times Flood = 1$ is whether the student lives in an area with a risk of flooding (defined as within a kilometer of the nearest body of water and with an address with an elevation below the 60 meter water mark) interacted with the alert of potential flooding (April 2017). This design is effectively a differences-in-differences, which compares students in low lying areas to their neighbors (the $Risk$ coefficient) in the post period (the $Flood$ coefficient).

In the first column, the $Risk \times Flood$ coefficient (functionally $Treat \times Post$) indicates a large and statistically significant depression of cognitive performance for students writing under the threat of flooding. This first column compares addresses below 60 meters to any of those above. Successive columns restrict the control group's elevation, to increase the comparability of the treatment and control. Smaller bandwidths approximate a regression discontinuity design, however it is unlikely that a sharp discontinuity exists as that would imply foreknowledge of the extent of the never before seen flooding levels. The second column limits the control group to address elevations up to 100m, the third column restricts further to 90m, and so on. As the window shrinks, the magnitude and significance of the estimated effect reduces only slightly until the final column, where the comparison is restricted to those living below 60m and those living just above up until 70 meters. While the reduced sample size could account for the slight increase in the standard errors, it is more probable that the anticipatory and depressive effect is not sharply identified at 60 meters. It is entirely possible that, if the mechanism is worries about being flooded, the lack of perfect forecasting could mean that students residing in marginal properties may also experience a depressive effect on their performance.

In Table 3, I conduct a similar exercise which analyzes the effects of the 2019 flooding. The analysis is restricted to students who were enrolled in both the 2018 and 2019 academic years, results are relatively similar to those seen in Table 2. A significant departure however is the amplified effect found in the smallest bandwidth for 2019, particularly when compared to the same bandwidth in 2017. It suggests that for those living in low lying areas and for those just above the water lines, those just above the water line do much better than those flooded. While at first glance this may seem to be a perverse result, it could be a manifestation of a near-miss phenom also notably highlighted in house prices in the United Kingdom that nearly miss being flooded (Beltrán et al., 2019). I study this question later in a section on non-linearity and present the results in Figures 7.

6.1.1 Heterogeneity by Student Characteristics

Lechowska (2018) reviews the literature on flood risk perceptions, and notes that researchers have found striking differences in the awareness and worry exhibited by affected populations by socioeconomic characteristics such as sex, immigration status and age. In Table 5, I combine the analysis of both floods and bisect the sample in a number of ways. Some interesting subsample results can be gleaned from the data, which may not be the focus of the current research, but may offer suggestive evidence towards some of the mechanisms at work.

First, I investigate possible gender differences in the threat to flooding. While the threat of flooding on female students has an estimated negative effect, it is not statistically significant. In contrast, male students living in low-lying areas are substantively and significantly affected. Gender differences in resiliency to environmental factors for exams are not new, for example in Cook and Heyes (2020), the authors noted that female students drawn from a similar population were less sensitive to extremely cold temperatures.

Second, I investigate whether foreign students, who are designated as such by their immigration status or having been required to pay 'international' rather than domestic student tuition, are differentially sensitive to flooding concerns as domestic students. A number of hypotheses could be put forward for possible reasons for the observed result that foreign students are not significantly perturbed by the threat of flooding - with both a larger standard deviation to the estimate and a much reduced estimated effect coefficient. For example, foreign students may be more likely to have experienced flooding or other disasters in their home countries. It may also be the case that foreign students are less likely to own (or be part of the family that owns) the address they are living in. I engage with the concept of renter and ownership later. For now, domestic students are significantly perturbed by the threat of flooding.

Third, I compare the estimated effect by student age group. The University (and others in the province) often treat students of 25 years of age differently than those younger, for example, through different admissions criteria. These 'mature' students are presumably more likely to own their homes, but a multitude of other differences likely exist. Whether someone in their mid-20's is differentially affected by the threat of flooding remains an empirical question. I find that students below the age of 25 are substantively perturbed by the threat of flooding as compared to mature students. Comparing this to the greater research body, the U.K. Environment Agency recently estimated that 18 to 34 year olds are less likely to know how to protect their homes, or where to go for flooding information.

6.1.2 Heterogeneity by Housing Characteristics

Lechowska (2018) reviews the literature on flood risk perceptions, and notes that researchers also found differences in the awareness and worry exhibited by affected populations by housing (residence) characteristics such as construction vintage and ownership. In Table 6, I combine the analysis of both floods and examine heterogeneity of the cognitive decrement by the characteristics of an addresses' housing. It is not unlikely that the characteristics of the built environment would have a defensive or protective effect, for example Beltrán et al. (2019) find flooding harmed lower-priced properties, and that apartments suffer less of a price discount than semi and detached housing. Bearing this in mind, I first consider if there are substantive differences between those living in houses or in taller apartments. Second, I consider if there are substantive differences between those living in areas with newer or older buildings. Third, I examine how high versus low rental costs are differentially connected to a productivity decrement.

In the first column, I restrict the sample to addresses that are in areas with no apartments greater than five stories. These areas are quite numerous in the setting, making up around half of the sample. These areas would represent students and workers that, should their area flood, are more likely to be personally affected. This is in comparison to the second column, wherein I restrict the sample to areas with a non-zero proportion of high rise apartment buildings. The first column shows a large and statistically significant effect of the threat of flooding on those living in lower buildings. Due to the sample restriction, the comparison here is students in low buildings at low elevations to students in low buildings at slightly higher elevations. When flooding is threatened, students in low elevation areas are substantially more affected. The second column shows a markedly smaller estimated coefficient effect size, resulting in an estimate that is statistically insignificant from zero. This lends evidence to the (perhaps) intuitive result that those living on higher floors of buildings may have a markedly different experience to flooding to those on the ground.

In the third and fourth column, I examine areas that have no new builds and areas with a non-zero amount of buildings completed between 2010 and 2016 (when the census was taken, and just before the flooding). When comparing mature neighborhoods to each other in the third column, it seems that one of two hypotheses could be at work. Either the effects of flooding are substantively different according to the relative age of the building (which could be driven by sump pumps or better city infrastructure for drainage) or, city planners may have recently begun to prohibit construction in flood prone areas. Regardless of which (or both) mechanism is at work, areas without any new building in recent years are substantively affected by the threat of flooding, while areas with a non-zero percent of homes built within the preceding 6 years are much less perturbed.

In the fifth and sixth columns, I connect the address of each student to the quality of housing, as proxied by the median rent in the dissemination area the students address is located. This is meant to test the possibility that lower quality housing is perhaps more perceived to be susceptible to flood damage than higher quality housing. I split the sample into above and below the median. Students in more expensive rental housing are less sensitive to the threat of the flooding as those in areas with lower costs to renting. Differential effects of flooding by house value on house transaction price were also found in a sample from the United Kingdom in Beltrán et al. (2019), whom also found wealthier areas to be more resilient.

6.1.3 Non-Linearity

In the current research setting, which examines whether there was a productivity decrement from the *threat* of a natural disaster, it is possible that using exposure data derived from after the disaster may be inappropriate. For example, using a sharp cut off at 60 meters above sea level as a binary indicator for potential flood exposure may ignore a productivity decrement of those marginal addresses just above the peak water level - and would not have known at the time they were not to be treated.

In Figure 7, I present the results of a non-linear regression of cognitive performance and flood risk. Instead of an indicator variable for whether an address is above or below the 60 meter water mark, I have instead grouped addresses into 10 meter elevation 'bins'. In particular, I use the non-linear differences-in-differences regression specification:

$$Grade_{i,e,t} = \alpha + \sum_{e=1}^{7} \gamma_e * Elevation_e + \lambda * Post_t + \sum_{e=1}^{7} \delta_e * (Elevation_e \times Post_t) + \nu_i + \varepsilon_{i,f,t} \quad (2)$$

Where $Grade_{i,e,t}$ is the imputed exam performance for individual i with an address in elevation interval e (for example between 60 and 70 meters is represented by $e = 3$), performing an exam in period t. My parameters of interest are then δ_e, the coefficients coinciding with the performance change of those warned of flooding, once the flood warning has been issued, binned at the 10 meter level. The inclusion of student (ν_i) fixed effects implies that identification continues to be derived from within-student variation. In other words, variations in the performance of individual subjects under alternative flooding treatments. Said differently, the interpretation of any of the δ_e is the average effect on grades of a student living in elevation e moving from the pre-control condition (e.g. 2016) to the treated condition (being warned of flooding, e.g. 2017).

The left panel of Figure 7 presents this analysis for the 2017 flooding. The leftmost coefficient (which is negative, large, and statistically significant even at the 5% level) indicates that students living in the lowest regions (anywhere from 40m to 50m above sea level) are most sensitive to the threat of flooding. While this would be consistent with the highest levels of damage and therefore anxiety (Lekuthai and Vongvisessomjai, 2001) it is also consistent with how the velocity (speed of onset) of flooding changes by elevation as well (Tingsanchali, 2012). The next coefficient represents the depressive effect of flood warnings on addresses between 50m and up to 60m. This critical region contains housing that can both reliably expect to be part of flooding, and others where only exceptional flooding with a '5%' chance of occurring would affect them. While statistically insignificant at the 5% level, the estimated coefficient is significant at the 10% level and retains a large negative magnitude. The reference group is addresses between 60m and 70m elevation, which, by construction have a zero estimated effect. Successive elevations do not have a statistically significant coefficient among them, but do exhibit a non negative trend to be expected as fluvial flooding is, without exaggeration, particularly unlikely to reach upwards of 100 meters.

The right panel of Figure 7 presents the non-linear analysis for the 2019 flooding. As in 2017 we see a particularly strong negative influence on the cognitive productivity of students living in the lowest two echelons of the city. The reference group once again has a mechanically zero estimated coefficient. Relative to the reference group however, is the "U" shaped curve following the reference point. This shape was hinted at in the final columns of Table 3, where the reference group seems to have done particularly well in comparison to both those most likely to have been affected by flooding and those above them. One explanation for this shape is the 'near miss' phenomenon which has been documented in other spaces, wherein those who had been *just* spared by flooding, may feel themselves to be irrationally protected from it.

6.1.4 Regression Discontinuity

In this section, I apply a regression discontinuity design to alternatively identify the effects of flood risk on cognitive performance.

Identification through more than one strategy demonstrates robustness to possible violations of the underlying assumptions required by a strategy. For example, if a violation of the parallel trends assumption in a differences-in-differences framework were responsible for producing the differences-in-differences results, then estimations of similar sign and magnitude using an alternative identification strategy (particularly one which does not rely on a time dimension) provides supporting evidence for the original framework. In my setting, I can use address elevation as the running variable for a discontinuity. Doing so implicitly states that students are only treated if they are below a known threshold. Consistent with the remainder of this paper, I choose this threshold to be 60 meters above sea level, just below the high water mark in either flood. What is notable about this discontinuity is that the running variable is obvious to the student - if not in absolute then in relative terms to the water. Further, students are unable to manipulate the elevation of their address (at least within a two week exam-period).

Specifically, I use a local linear regression on either side of the threshold, capturing the non-linearity around the kink point and the roughly linear effects estimated on either side. The specification is two staged. First, I remove both student and semester averages from the cognitive performance using ordinary least squares:

$$Grade_{i,s} = \alpha + \nu_i + \sum_{s=1}^{10} \gamma_s + r_{i,s} \tag{3}$$

Where $Grade_{i,s}$ is the grade achieved by student i in semester s. I keep 10 winter semesters for the analysis, from Winter 2010 through Winter 2019, inclusive. I then take residuals, $r_{i,s}$ and use them as the dependent variable in the discontinuity specification (separately by year):[20]

$$r_i = \beta_0 + \tau D + \beta_1(e_i - 60) + \beta_2 D(e_i - 60) + \varepsilon_i \tag{4}$$

Wherein the residuals are regressed against a constant, an indicator variable that takes the value of one if the address has an elevation above 60 meters, and linear coefficients on either side of the threshold. Specifically, τ represents the average magnitude

[20] This non-parametric method (the local linear regression) should help to reduce the bias that can result from data further from the cutoff, specifically, from higher elevations. For example fitting a quadratic or cubic polynomial to the 2017 and 2019 data could result in higher elevations significantly bending the polynomial near the cutoff, an undesirable result in this setting.

(up to weighting) of the discontinuity, β_1 estimates the slope of the performance-elevation gradient below the threshold and β_2 estimates the slope of the performance-elevation gradient above the threshold. I employ the common triangular kernel for weighting purposes. I present the results in Figure 9 and also in Table 9.

In Figure 9, I present four discontinuity plots.[21] The top left panel depicts results for 2016 where no exceptional floods have taken place, either during the exam period nor in recent memory (see Figure 1). Each marker represents the average performance deviation for addresses with elevations within each bin. The performance deviation, from the two-step procedure, is the deviation after within-student averages and all student semester averages have been removed. A vertical line is depicted at the threshold of 60 meters, coinciding closely with the water levels not seen before but within the realm of possibility. Linear regressions, fitted on either side of the threshold, are depicted with solid lines. In 2016, the below-threshold performance-elevation gradient is nearly flat, almost consistent with the above threshold gradient. The vertical difference between where the two estimated gradients meet is small, indeed the discontinuity estimate is small, at an average of 0.281 percentage points, with an associated t statistics of only 0.52.

The top right panel depicts results for 2017, where minor flooding was occurring during the exam period and exceptional flooding was forecast as possible. Here, the above threshold performance gradient remains near flat. However, the before threshold performance elevation gradient is strongly negative (that is, more marginal properties have a higher productivity decrement than less marginal ones). The vertical difference between those above and below the threshold is statistically significant - with a difference of 0.896 percentage points and an associated t statistic of 1.71 (statistically significant at the 10% level).

[21] I use these plots as the primary presentation of my RD results as plots such as these are widely called for in RD designs and also play a role in their validation (Cook, 2008; Imbens and Lemieux, 2008; Lee and Lemieux, 2010).

The bottom left panel depicts results for 2018, where no flooding was forecast and indeed, *laissez-faire* press releases were given by the regulator. The estimated shape is relatively similar to that seen in 2016, with a relatively flat above- and below-threshold performance gradient. A slight difference however, is the vertical difference between the two populations; those below the high water mark threshold in 2018 perform *better*.[22]

The bottom right panel depicts results for 2019, where the worst flooding would later occur and more significant flood warnings were issued. It is also possible that the city's (and even maybe the students') previous experiences with flooding could affect the relationship between elevation and performance. Indeed, the before-threshold elevation-performance gradient is estimated to be positive, a shift from the negative estimates in all preceding years. This indicates that more marginal properties have a less severe productivity decrement than those at the lowest elevations, consistent with correctly updating beliefs from recent flooding. The vertical difference between those above and below the threshold is once again statistically significant - with a difference of 0.918 percentage points and an associated t statistics of 1.95 (statistically significant at the 10% level).

6.1.5 Within-Student Within-Season

In this section, I use the day an exam is written to leverage *within-student* and *within-exam-period* variation for identification. Specifically, I use the equation:

$$Grade_{i,t} = \alpha_0 + \alpha_1 \mathrm{X_{i,t}} + \nu_i + \varepsilon_{i,t} \tag{5}$$

Where $Grade_{i,t}$ is the grade achieved in a course by student i writing an exam on day t. I separately estimate the equation for every exam season (in other words 2016 is conducted separately from 2017 and so on) allowing, for example, the α_0 terms to

[22] This could be due to selection effects - students affected in 2017 are more likely to subsequently drop out of studies and defensively emigrated from low lying areas - see the following sections on persistence and migration.

vary by year. The $X_{i,t}$ term indicates a treatment variable where identification will come from 'switchers', those whose status moves from untreated to treated. In this way, students act as their own control groups. For example, in Table 11, the treatment variable changes from zero to one when the measured water level is higher than student i's address elevation. The student fixed effects (ν_i) remove the influence of student-level time-invariant unobserved characteristics, where much should be time invariant as the exam seasons are typically two weeks.

In Table 10, I examine if precipitation has a differential effect on performance depending on student address elevation. My findings coincide with the cognitive effects of rain and flooding found elsewhere in the literature; Decent (2018) found long-term adverse mental-health effects from rain in households that had been affected by flooding as much as three years after the event. They found that nearly half of homeowners who had experienced flooding report increased anxiety whenever it rains compared to three per cent of homeowners in the 1 kilometer surrounding control group. Daily precipitation data was collected from the nearby weather monitor.[23] I find that total daily precipitation is not usually associated with either an increase or decrease in performance *unless* a student lives in a low lying area, when precipitation has a consistently estimated negative effect. This precipitation effect is only statistically significantly different from the precipitation effect of those who live on higher ground during exam seasons when flooding would eventually occur. Because precipitation is also substantially higher in exam seasons with eventual flooding, I caution over-interpretation of this result as it may be a response to dosage rather than the ability of those in low-lying areas successfully predicting flooding from the rain. Nonetheless, a within-student one standard deviation increase in precipitation of just over four millimeters coincides with a reduction of almost 4 percent of a standard deviation in 2017 and in 2019, a one stan-

[23]While snow is unlikely in April, it has been known to occur. For this reason I use total precipitation which converts snow into it's equivalent water volume at a ratio of approximately ten-to-one in volume: If one centimeter of snow were to fall, total precipitation would record a total of one millimeter of precipitation.

dard deviation increase in precipitation of over 9mm causes a reduction in performance by more than five percent of a standard deviation.

In Table 11, I examine the effect of student performance when their address is most likely flooded. This exercise most closely coincides with the literature of the *direct* effects of exposure to natural disasters. To do this, I compare the elevation of a student's address to the daily maximum water level of the river. If the daily maximum water level is above the student's address, I use an indicator variable value of one to assign the treatment. As it is a panel-fixed effect model, the coefficient on a binary treatment variable is identified only on 'switchers', or students whose addresses change from being in an unflooded to the flooded state. Accordingly, the bottom panel of Table 11 notes the number of students that, from the population of those living within 1 kilometer of the river and below 100 meters, live low enough to have at least one of their exams written when exposed to flooding. An estimated three to four hundred students are affected in this way, with more affected with higher maximum water levels. The lowest number affected is in 2018, both when water levels were on average the lowest and immediately after a 100-year flood. Fluvial flooding is common - the threshold for a minor, common, or '2-year' flood is 59.42 meters. This means that a minor (or greater flood) occurred *during* the exam seasons of 2016, 2017, and 2019 (I provide both the average water level and the maximum during the exam season in the bottom panel of Table 11). I graphically present the evolution of the water level and the exam period in Figure 3 for 2017 and in Figure 4 for 2019. I estimate the equation separately by year. It is perhaps unsurprising that a relatively large performance decrement is identified here. Within-student and within-exam season, I find that when exposed to flooding students perform around *half of a standard deviation worse* than during exams when they are not exposed to flooding.

6.1.6 Website Traffic

In Table 12, I present the (estimated as linear) relationship between the river's water level and the number of pageviews at the Regulator's flood warning website.[24] My prior is that as water levels rise, so too would the worries about them. While there does not exist readily available before, during, and after anxiety levels, I can use access statistics shared by the primary flood warning website `www.ottawariver.ca` as a proxy for how worried people were.

The results confirm that as water levels rise so do concerns about flooding. In 2016, during a once-every-two-year flood, a one centimeter increase in water level had an estimated effect of an additional 15 pageviews. During the beginnings of a once-in-a-hundred year flood for 2017, a one centimeter increase in water level had an estimated effect of an additional 70 pageviews. In 2018, a year with comparatively unremarkable water levels a one centimeter increase in water level was (not statistically significantly) associated with an additional 3 pageviews. During the second once-in-a-hundred year flood in the time period, a one centimeter increase in water level increased pageviews by 228. In Figure 3 and Figure 4 I have also included the evolution of the pageviews.

In Table 14, I leverage the pageviews as a proxy for worries about flooding to estimate the effects of those worries on performance. In 2016, there is no statistically significant relationship between the number of pageviews and student performance either below 60 meters or above. However, both are negative, suggesting that more pageviews are associated with less performance. In 2017, when a once-in-a-hundred year flood would occur, for students above 60 meter threshold, their performance *increased* with increasing levels of accessing the flood warnings website. Conversely, those living in low-lying areas were had *decreased* performance with increases in accessing the warnings site. A one standard deviation increase in the number of pageviews had an estimated effect of widening the gap between the two populations by almost four per-

[24] In Table 13, I present the same using unique 'new visitors'.

cent of a standard deviation. The relationship returns to non-statistical significance in 2018, a year with negligible water levels. In 2019, when the second once-in-a-hundred year would occur, a one standard deviation increase in pageviews was again associated with a 4.7% of a standard deviation decrease in performance for those in low lying areas.

6.1.7 Direct Effects of Warnings

In this section, I examine the direct effects of flood warnings by comparing the exams written by a student before and after a major flood warning was issued. There is a notable difference between the two years to keep in mind. For 2017, students began exams under a warning of minor (routine) flooding, whereas for 2019 the first major warning came well into the exam season. To identify the marginal effects of a major flood warning issue (for both years, the major flood warning is made some time after a routine flood warning) I use Equation 5 with the treatment variable assigned to exams written on the day of, or after, the major flood warning was issued. I present the results in Table 15.

In the first column, I estimate the effect of a major flood warning within the 2017 exam season. Briefly summarizing the content of Section 3.1, a minor flooding alert was in issue before exams even began in 2017. About halfway, an alert was issued that major flooding could be expected. I present this timeline graphically in Figure 3. I estimate the effect of a flood warning with a binary indicator variable that takes the value of one on April 18, 2017 (the day the alert was issued) and for any day afterwards. All students, regardless of their elevation suffer a 4% of a standard deviation reduction in performance upon the release of the news. Those who would presumably be at risk at addresses below 60 meters are also negatively affected (although the standard errors remain too large for statistical significance, the effect is of a congruent size and magnitude).

In the second column, I estimate the effect of a flood warning within the 2019 exam season. Briefly summarizing the content of Section 3.2, there was no flood issue leading into the 2019 exam season. On April 16, a warning was issued that water levels may be minor, with levels below that seen two years prior. On April 18, a major flood warning that levels exceeding the previous flood could be expected. I present this timeline graphically in Figure 4. I estimate the effect of a flood warning with a binary indicator variable that takes the value of one on April 18, 2019 (the day the alert was issued) and for any day afterwards. The sensitivity to this news for all students is negligible; only students who are at or below the 2017 water mark are negatively affected. The estimate of the effect is over 9% of a standard deviation, larger than the effect from 2017.

6.2 Persistence and Government Disaster Recovery Assistance

Another outcome variable, with arguably more serious repercussions than short-term cognitive performance reductions, is stopping studies altogether without completing them. Connecting short term or transitory shocks to long run outcomes is an essential part in correctly estimating the magnitude of the shocks. It is also the defining feature of recent causal analysis of environmental hazards. For example, Zivin and Neidell (2009) connects short term air pollution to the ability to attend university. Park (2020) connects exam day temperatures to the likelihood of eventual high school graduation. In this section, I provide evidence that exposure to flooding results in a large, statistically significant increase in the probability of dropping out - ceasing studies without graduation. I also show that eligibility for government disaster relief alleviates this effect.[25]

In both panels of Figure 8, I analyze how students living in low-lying areas dif-

[25] Disaster Recovery Assistance for Ontarians, administered by the Ontario Provincial Government. In Figure 2, the red line along Carling Avenue also represents the border of DRAO eligibility for that area. Other polygons are presented in Figures

ferentially cease studies compared to their higher-ground neighbors. For this exercise, I compare the four semesters before and four semesters after the May 2017 flooding (indicated with a vertical line between Winter 2017 and Fall 2017). The horizontal axis from left to right indicates the passage of time, including only fall and winter semesters.[26] The vertical axis indicates the magnitude of the underlying regression coefficient. The regression's dependent variable is a binary (indicator) variable that takes the value of one if a student was no longer taking classes and had not graduated in that semester. The unit of analysis is the student-semester. The fully specified differences-in-differences regression equation is:

$$Dropout_{i,t} = \alpha + \sum_{t=1}^{T} \gamma_t * Semester_t + \lambda * Treated_i + \sum_{t=1}^{T} \delta_t * (Semester_t \times Treated_i) + \varepsilon_{i,t} \tag{6}$$

Where the coefficients depicted in Figure 8 are the coefficients of the interaction terms, δ_t.

In the left panel of Figure 8, I estimate Equation 6 for students in areas that are not eligible for disaster relief under the DRAO. For the four semesters prior to the 2017 flooding, students in low-lying areas are no more or less likely to drop out of classes. For the first semester following the floods (Fall 2017), students at lower elevations are more than 4 percentage points less likely to return to studies. In the next semester, Winter 2018, the magnitude and statistical significance of this excess dropout are relatively unchanged, consistent with the intuition that the option to dropout is exercised at the beginning of an academic year (rather than partway through). This effect persists, albeit with a slightly lower estimate in the 2019 academic year as well. To fix ideas about the magnitude of these estimated effects, an average of 13% student-semesters

[26]Summer study is available at the university, however it is not common nor customary, as it is when only 6% of class grades are achieved.

are classified as 'drop-out' during this period.[27]

In the right panel of Figure 8, I estimate Equation 6 for students in areas that are eligible for disaster relief under the DRAO. I present reconstructions of the five distinct polygon areas in Figure A1. When activated, homeowners are eligible to apply for provincial government assistance under the "Disaster Recovery Assistance for Ontarians" program (more information here and here). In a sudden natural event such as a flood, the DRAO reimburses expenses to homeowners for clean up, repair and replacement of essentials such as furnaces and hot water heaters, and basic emergency expenses like travel evacuation costs. Damage caused by overland flooding, infiltration flooding and sump pump failure is eligible for assistance under the program. There is a \$250,000 per application limit, a \$500 waivable deductible, and a co-pay of 10%. For the four semesters prior to the 2017 flooding, students in low-lying areas that were eventually covered by the DRAO are no more or less likely to drop out of classes. For the first and second semester following the floods (Fall 2017 and Winter 2018), students at lower elevations yet covered by the DRAO (which comprise around one-tenth of the students in question) are not more or less likely to drop out than their higher ground peers. In fact, the estimates for this first year are negative, although statistically insignificant. The magnitude within an academic year is relatively equal, again consistent with the intuition that the option to dropout is exercised at the beginning of an academic year (rather than partway through). The following year, those students in the low-lying sections of DRAO relief zones are estimated to be *less* likely to drop out than those in DRAO zones at higher elevations.

[27] The effect of the 2019 flooding, which occurred during and after the Winter 2019 semester cannot be measured with the current data which ends in Winter 2019. Even as the data becomes available, it is unclear how the multitudinous and idiosyncratic effects of COVID-19 should be accounted for.

6.3 Migration and Displacement

In this section, I analyze how students change their behavior after being exposed to the risks of flooding. Prior to 2017, many areas that flooded were not considered at risk, or if they were, that risk was not salient. After the flooding, I find that students living close to the water or at low elevations persistently move further from the water and to higher ground.

In Table 7, I analyze how students exposed to the 2017 flooding changed address from year to year. In all columns, the dependent variable is the distance of a student's address from the nearest body of water, measured in meters. In the first column, I compare the average distance to water change from 2015 to 2016, a year before the flooding occurred. Students who would not later be exposed to the flooding moved closer to the water an average of 8 meters, but the estimate is far from statistically significant. In contrast, those who would later be exposed moved an average of 67 meters closer to the water. This result is expected, as almost mechanically, students needed to move closer to the water to be threatened by the later 2017 flooding. Nonetheless, the trend continues when comparing spring 2016 and 2017. The striking result is the reversal of the trend beginning the year after the flooding. In the third column, those exposed to the 2017 flood warnings move an average of 78.5 meters *further* from the water. This is contrast to those students who were not exposed, who move closer to the water on average. The trend continues in the fourth column comparing 2018 addresses (no flooding) to 2019 addresses (just before the second flood). I identify that students who were exposed to the first flooding continue to move further away from the water, *even compared to* the 2018 addresses that were already distanced.

In Table 8, I find similar results; before the flooding a move to lower elevations and afterwards exposed students have a distinct and persistent movement to higher ground.

7 Conclusion

Climate change has seen an increase in both the incidence and severity of natural disasters. Fires, floods, and other environmental hazards are becoming the new normal in places that have not historically been afflicted by them, and in places disasters normally strike, they are becoming more costly and deadly.

In a North American setting which experiences regular minor flooding relating to the spring thaw (or freshet), two 100-year floods occurred with only one year in between. With warning systems in place and without significant ability to directly control the volumes of water, this double-shock allows for the identification of a large and statistically significant productivity decrement when students are warned their homes could be flooded.

I find that when the threat of a flood looms large, students below the likely water line are around 7% of a standard deviation *less* capable of completing strenuous cognitive tasks, despite being well incentivized. The direction and order of magnitude are similar when using either a differences-in-differences identification strategy or one based on a regression discontinuity design.

I find evidence that certain adaptations are effective in reducing the effects of flood warnings. For example, structures such as taller buildings have a neutralizing effect on the estimated productivity decrement, suggesting that the built environment could play a key role in protecting students and workers at large. I also find that living in newer structures and living in buildings with comparatively high rents[28] (which serves as a proxy for housing *quality*) were also successful in mitigating the anticipatory productivity decrement from flooding.

I also find evidence that students adapt through migration. Identifying those that were likely exposed to flooding, I find that they relocate to addresses both further

[28] Higher rents before the flooding, it is possible rents changed after a building is, and is known to be, flood prone. An area for future research.

from the nearest source of fluvial flooding (by around 80 meters on average) and, most importantly, to higher elevations (around 2.3 meters on average, roughly the extent to which the fluvial flooding was 'above norm').

Lastly, I identified long term effects on those treated; when students are treated to flooding they are almost one third more likely to drop out of studies altogether the following year. While students could return to school after the vintage of my data set is sufficient, a year long interruption in human capital accumulation or even in life-cycle earnings can have long term consequences. If the identified effects herein are sustained in the greater economy, the forecasted increase in the incidence of natural disasters may have an additional and currently unaccounted for affect on the economy.[29]

[29] Hsiang and Jina (2014) provide an important context when examining environmental impacts on the economy; small estimated effects (in their case from additional tropical cyclone exposure) can have significant and large long term effects 'not obvious to the casual observer'. My paper has attempted to causally identify the effects using a particularly sensitive cognitive productivity measure. While it remains to be seen if these floods will have a long term effect (the most recent occurred less than two years ago) it is notable that something such as dropping out of studies can have persistent effects on lifetime earnings for the individual and their economy.

References

Angrist, J. D. and Pischke, J.-S. (2008). *Mostly harmless econometrics: An empiricist's companion*. Princeton University Press.

Archsmith, J., Heyes, A., and Saberian, S. (2018). Air quality and error quantity: Pollution and performance in a high-skilled, quality-focused occupation. *Journal of the Association of Environmental and Resource Economists*, 5(4):827–863.

Barwick, P. J., Li, S., Lin, L., and Zou, E. (2019). From fog to smog: The value of pollution information. Technical report, National Bureau of Economic Research.

Beatty, T. K., Shimshack, J. P., and Volpe, R. J. (2019). Disaster preparedness and disaster response: Evidence from sales of emergency supplies before and after hurricanes. *Journal of the Association of Environmental and Resource Economists*, 6(4):633–668.

Belasen, A. R. and Polachek, S. W. (2008). How hurricanes affect wages and employment in local labor markets. *American Economic Review*, 98(2):49–53.

Beltrán, A., Maddison, D., and Elliott, R. (2019). The impact of flooding on property prices: A repeat-sales approach. *Journal of Environmental Economics and Management*, 95:62–86.

Cook, N. and Heyes, A. (2020). Brain freeze: Outdoor cold and indoor cognitive performance. *Journal of Environmental Economics and Management*, page 102318.

Cook, T. D. (2008). "Waiting for life to arrive": A history of the regression-discontinuity design in psychology, statistics and economics. *Journal of Econometrics*, 142(2):636–654.

Dassanayake, D. R., Burzel, A., and Oumeraci, H. (2015). Methods for the evaluation of intangible flood losses and their integration in flood risk analysis. *Coastal Engineering Journal*, 57(1):1540007–1.

Davies, J. B. (2020). Reforming Canada's disaster assistance programs. *Canadian Public Policy*, 46(2):187–197.

Decent, D. (2018). *After the Flood: The impact of climate change on mental health and lost time from work*. Intact Centre on Climate Adaptation.

Dell, M., Jones, B. F., and Olken, B. A. (2012). Temperature shocks and economic growth: Evidence from the last half century. *American Economic Journal: Macroeconomics*, 4(3):66–95.

Ebenstein, A., Lavy, V., and Roth, S. (2016). The long-run economic consequences of high-stakes examinations: Evidence from transitory variation in pollution. *American Economic Journal: Applied Economics*, 8(4):36–65.

Field, C. B. (2014). *Climate change 2014–Impacts, adaptation and vulnerability: Regional aspects*. Cambridge University Press.

Graff Zivin, J. and Neidell, M. (2014). Temperature and the allocation of time: Implications for climate change. *Journal of Labor Economics*, 32(1):1–26.

Groen, J. A. and Polivka, A. E. (2008). The effect of Hurricane Katrina on the labor market outcomes of evacuees. *American Economic Review*, 98(2):43–48.

Harrison, G. W. and List, J. A. (2004). Field experiments. *Journal of Economic literature*, 42(4):1009–1055.

Henstra, D. and Thistlethwaite, J. (2017). *Climate change, floods, and municipal risk sharing in Canada*. Institute on Municipal Finance and Governance.

Hossain, F. (2019). Creative destruction or just destruction? Effects of floods on manufacturing establishments in India. Job Market Paper - University of Melbourne.

Hsiang, S. and Kopp, R. E. (2018). An economist's guide to climate change science. *Journal of Economic Perspectives*, 32(4):3–32.

Hsiang, S. M. (2010). Temperatures and cyclones strongly associated with economic production in the Caribbean and Central America. *Proceedings of the National Academy of Sciences*, 107(35):15367–15372.

Hsiang, S. M. and Jina, A. S. (2014). The causal effect of environmental catastrophe on long-run economic growth: Evidence from 6,700 cyclones. Technical report, National Bureau of Economic Research.

Imbens, G. W. and Lemieux, T. (2008). Regression discontinuity designs: A guide to practice. *Journal of Econometrics*, 142(2):615–635.

Keiser, D., Lade, G., and Rudik, I. (2018). Air pollution and visitation at us national parks. *Science Advances*, 4(7):1613.

Kulp, S. A. and Strauss, B. H. (2019). New elevation data triple estimates of global vulnerability to sea-level rise and coastal flooding. *Nature Communications*, 10(1):1–12.

Lechowska, E. (2018). What determines flood risk perception? a review of factors of flood risk perception and relations between its basic elements. *Natural Hazards*, 94(3):1341–1366.

Lee, D. S. and Lemieux, T. (2010). Regression discontinuity designs in economics. *Journal of Economic Literature*, 48(2):281–355.

Lekuthai, A. and Vongvisessomjai, S. (2001). Intangible flood damage quantification. *Water Resources Management*, 15(5):343–362.

McNeil, D. (2019). An independent review of the 2019 flood events in Ontario.

Meier, K. J., O'Toole Jr, L. J., and Hicklin, A. (2010). I've seen fire and I've seen rain: Public management and performance after a natural disaster. *Administration & Society*, 41(8):979–1003.

Moretti, E. and Neidell, M. (2011). Pollution, health, and avoidance behavior evidence from the ports of Los Angeles. *Journal of Human Resources*, 46(1):154–175.

Ottawa River Regulation Planning Board (2017). Summary of the 2017 spring flood. Technical report.

Park, R. J. (2020). Hot temperature and high stakes performance. *Journal of Human Resources*.

Parker, D. J. (2017). Flood warning systems and their performance. In *Oxford Research Encyclopedia of Natural Hazard Science*.

Rabassa, M. J., Grand, M. C., García-Witulski, C. M., et al. (2020). Heat warnings and avoidance behavior: Evidence from a bike-sharing system. *Environmental Economics and Policy Studies*, pages 1–28.

Saberian, S., Heyes, A., and Rivers, N. (2017). Alerts work! Air quality warnings and cycling. *Resource and Energy Economics*, 49:165–185.

Scott, B. G., Lapré, G. E., Marsee, M. A., and Weems, C. F. (2014). Aggressive behavior and its associations with posttraumatic stress and academic achievement following a natural disaster. *Journal of Clinical Child & Adolescent Psychology*, 43(1):43–50.

Strömberg, D. (2007). Natural disasters, economic development, and humanitarian aid. *Journal of Economic Perspectives*, 21(3):199–222.

Tingsanchali, T. (2012). Urban flood disaster management. *Procedia engineering*, 32:25–37.

Yu, C.-H., Zhu, J.-D., and McCarl, B. A. (2019). Market response to typhoons: The role of information and expectation. Technical report, Agricultural and Applied Economics Association.

Zhang, P., Deschenes, O., Meng, K., and Zhang, J. (2018). Temperature effects on productivity and factor reallocation: Evidence from a half million Chinese manufacturing plants. *Journal of Environmental Economics and Management*, 88:1–17.

Zivin, J. G. and Neidell, M. (2009). Days of haze: Environmental information disclosure and intertemporal avoidance behavior. *Journal of Environmental Economics and Management*, 58(2):119–128.

Zivin, J. G., Song, Y., Tang, Q., and Zhang, P. (2020). Temperature and high-stakes cognitive performance: Evidence from the National College Entrance Examination in China. *Journal of Environmental Economics and Management*, 104:102365.

8 Figures

Figure 1: Exceptional Water Levels in 2017 and 2019

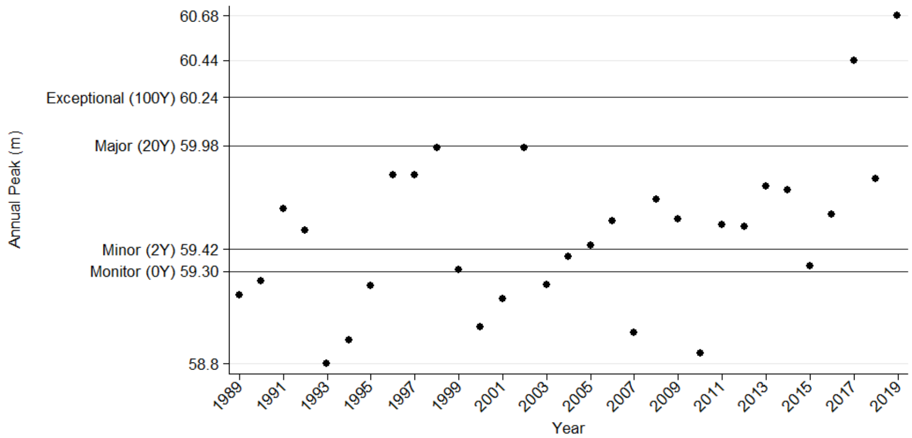

Exceptional water levels during the 2017 and 2019 flooding. Thirty years of data shown. The vertical axis is the annual peak water level (regardless of month achieved). When the water level exceeds 59.30 m, the regulator's status is set to 'monitor'. When the water level exceeds 59.42 m, the status is set to '2 Year Flood'. When the water level exceeds 59.98 m, the status is set to '20 Year Flood'. When the water level exceeds 60.24 m, the status is set to '100 Year Flood'. Data from the Ottawa River Regulation Planning Board.

Figure 2: Density of Postal Codes

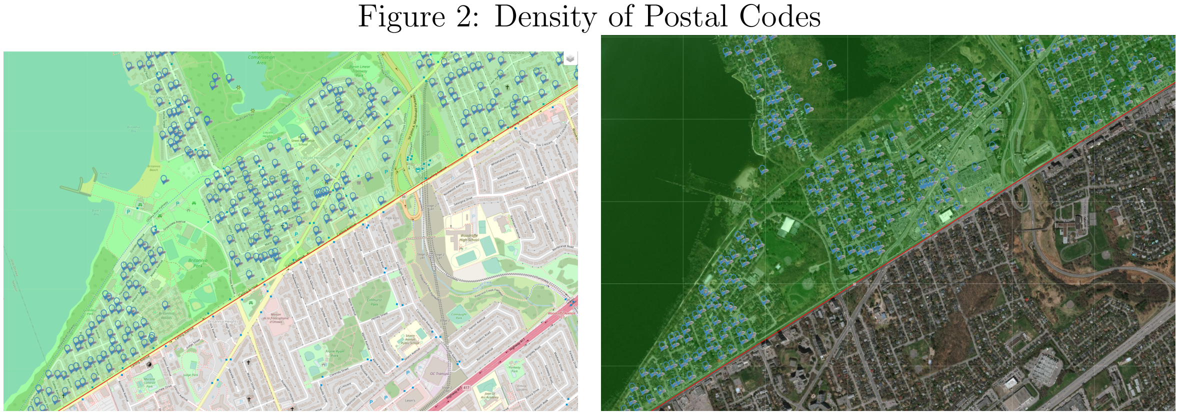

The density of six-character postal codes in a section of Ottawa's West End. Each point represents a unique six-character postal code. Decimal coordinates for the beach pictured are (45.36545, -75.80151)

Figure 3: Water Levels and Flood Warnings 2017

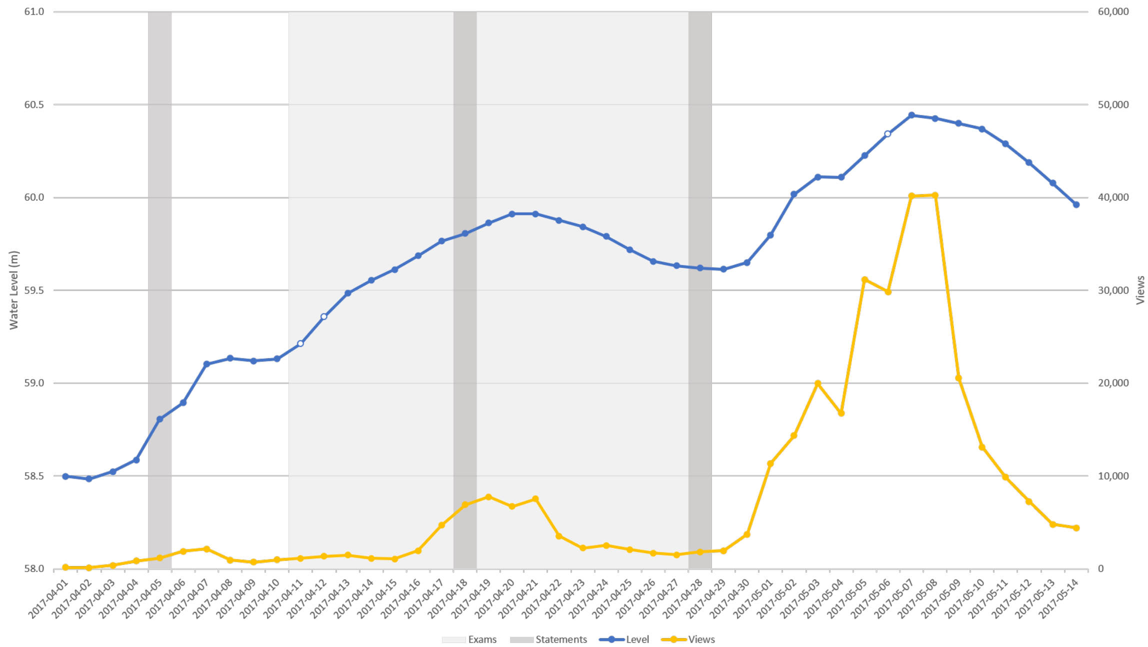

Light shading indicates the exam period. Dark shading indicates days when a flooding press release was issued. Blue (dark in print) indicates measured water level in meters. Yellow (light in print) indicates the number of pageviews the flood warning website received. Hollow points indicate a change in flood status. When the water level exceeds 59.30 m, the status is set to 'monitor'. When the water level exceeds 59.42 m, the status is set to '2 Year Flood'. When the water level exceeds 59.98 m, the status is set to '20 Year Flood'. When the water level exceeds 60.24 m, the status is set to '100 Year Flood'.

Figure 4: Water Levels and Flood Warnings 2019

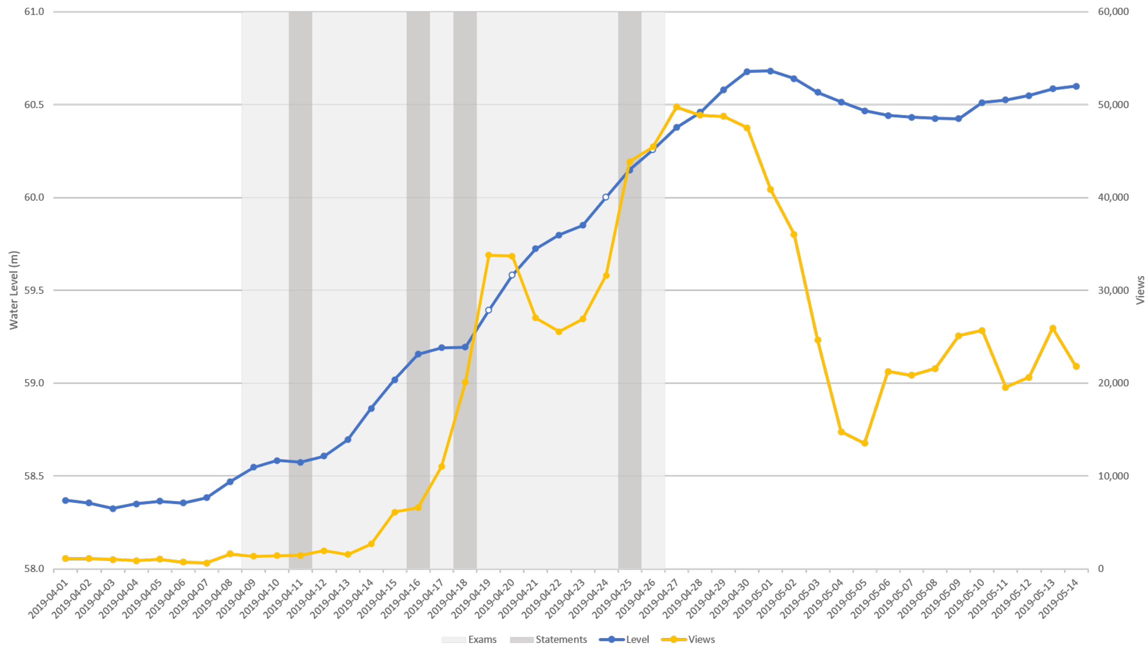

Light shading indicates the exam period. Dark shading indicates days when a flooding press release was issued. Blue (dark in print) indicates measured water level in meters. Yellow (light in print) indicates the number of pageviews the flood warning website received. Hollow points indicate a change in flood status. When the water level exceeds 59.30 m, the status is set to ‘monitor’. When the water level exceeds 59.42 m, the status is set to ‘2 Year Flood’. When the water level exceeds 59.98 m, the status is set to ‘20 Year Flood’. When the water level exceeds 60.24 m, the status is set to ‘100 Year Flood’.

Figure 5: Elevation of Postal Codes

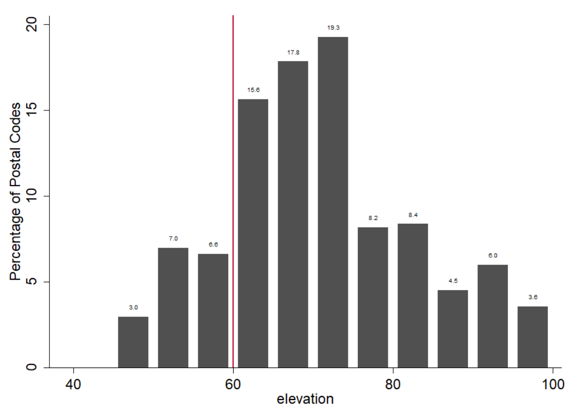

Histogram of address elevations. Addresses within 20km of campus used. Addresses within 1km of water used. Bars are indicated with the percent of postal codes they represent.

Figure 6: Distance to River

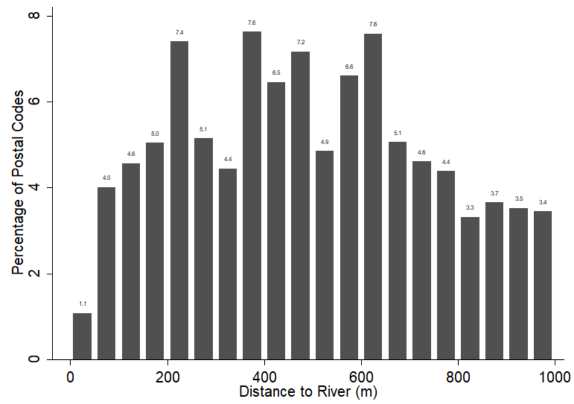

Histogram of distance to nearest river, in meters. Addresses within 20km of campus used. Addresses within 1km of water used. Bars are indicated with the percent of postal codes they represent.

Figure 7: Non-Linearity of Performance Decrement in 2017 and 2019

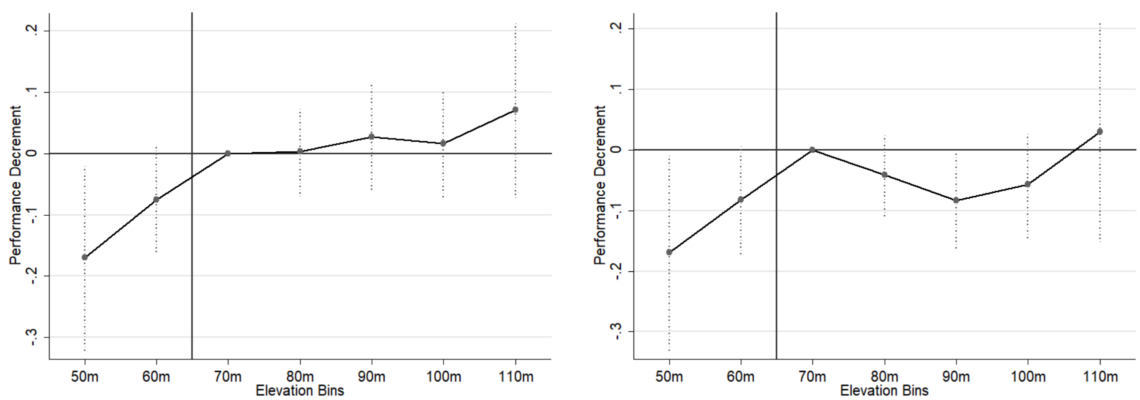

Regression coefficients for the interaction term of elevation and treatment. The left panel presents coefficients for 2017. The right panel presents coefficients for 2019. The dependent variable is standard deviation in exam score. The regression also contains indicator variables for each elevation and for the post period. Vertical reference line provided at the elevation the flooding did not exceed. Within-student fixed effects model. 95% confidence whiskers included. See Equation 2.

Figure 8: Excess Dropout Rate By Disaster Recovery Status

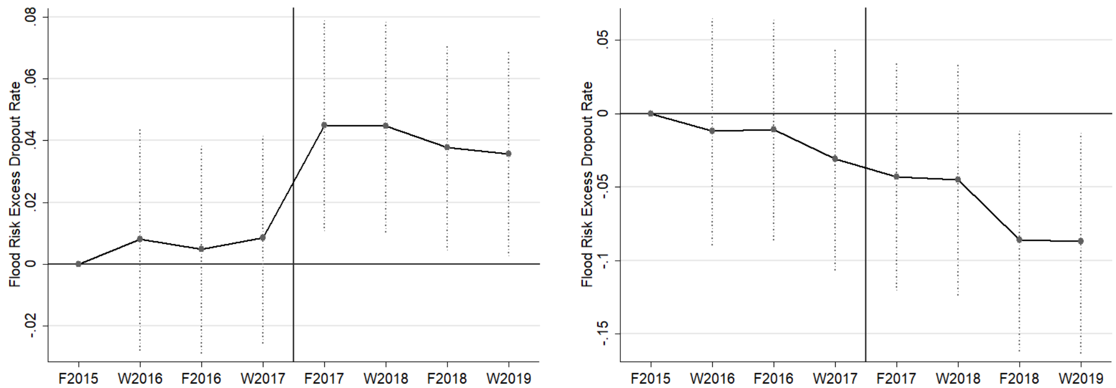

The left panel is restricted to students not covered by the 2017 Disaster Recovery Assistance for Ontarians activation. The right panel displays students covered by the DRAO activation in 2017. Presented are regression coefficients of the interaction between living below the 60m flood mark and semester. The dependent variable is a binary variable that equals one if a student-semester represents a drop out (taking no courses after having taken courses the semester previous without later graduation). The regression also contains dummy variables for each semester. Vertical reference line provided when flooding occurred in 2017. 13% of student-semesters during the time frame are considered dropouts. 95% confidence whiskers included. See Equation 6.

Figure 9: Regression Discontinuity Approach

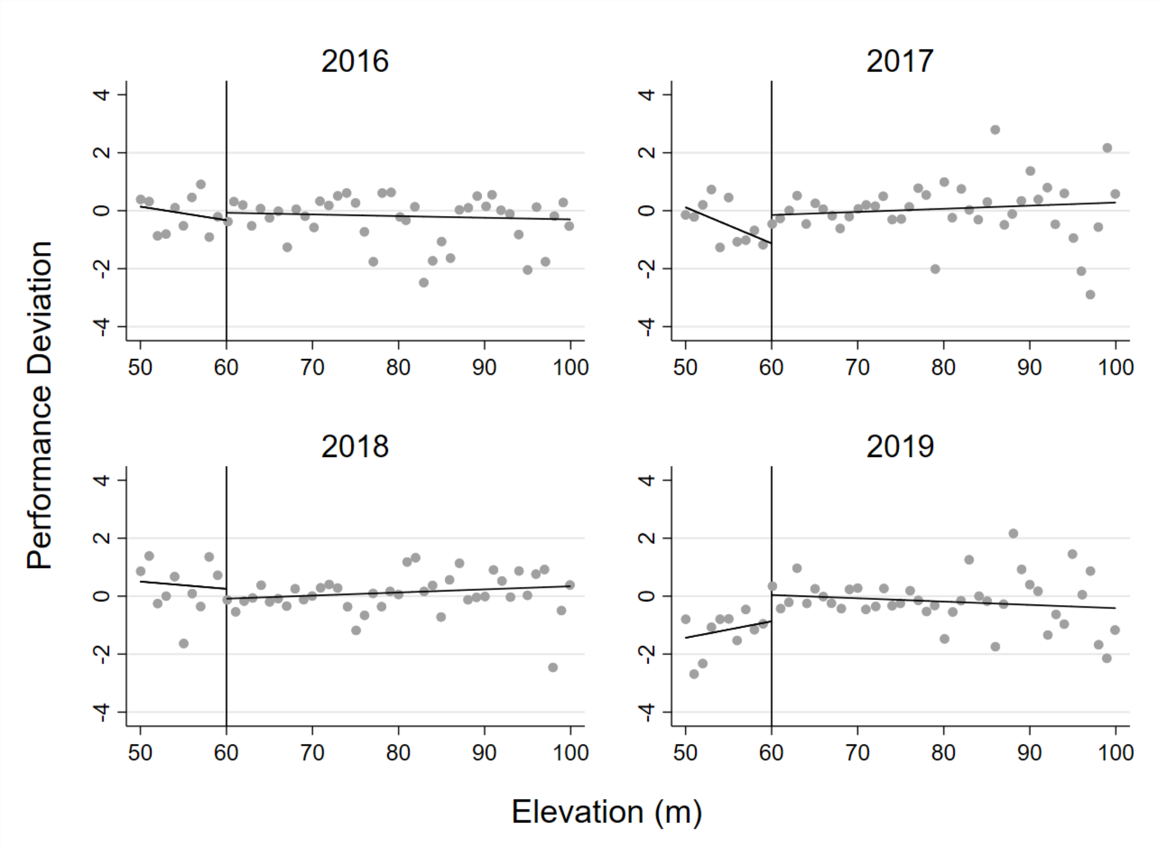

Regression discontinuity plots for each year in the sample separately. The horizontal axis is postal code elevation above sea level. The vertical axis is performance deviation in percent. Markers represent the average performance deviation within their elevation bin. Lines indicate estimated linear functions. Using a support of 10 meter radius, the 2016 discontinuity is 0.281 (t = 0.52), 2017 is 0.896 (t = 1.71), 2018 is -0.450 ($t = -0.97$), and 2019 is 0.918 ($t = 1.95$) Also see Table 9.

9 Tables

Table 1: Summary Statistics

	Mean	Std. Dev.
Outcomes		
Percent Grade	75.605	14.289
Spatial Characteristics		
Elevation (m)	70.960	14.458
Distance to Water (m)	474.328	244.275
Distance to Campus (km)	5.408	5.851
Student Characteristics		
Full Time	0.893	0.309
Age At Exam	21.696	4.122
Female	0.567	0.496
Foreign	0.267	0.442
Median Rent of Area	979.497	513.805
Observations	67,656	
Students	9,308	
Postal Codes	3,159	

Author's calculations. Also discussed in Section 4. Summary statistics for students with course grades completed during the Winter 2016, Winter 2017, Winter 2018, and Winter 2019 semesters. Students whose address is within 1,000 meters of major water source. Students whose address is at or below 110 meters above sea level. Students with addresses within 20 kilometers of campus. Full Time is defined by the University as taking 3 or more courses per semester. Median Rent from the 2016 Canadian Census of Population.

Table 2: The Effect of 2017 Flooding Risk on Performance

	(1) Score	(2) Score	(3) Score	(4) Score	(5) Score
Risk=1 × Flood=1	-0.073**	-0.072**	-0.071**	-0.067*	-0.057
	(0.035)	(0.035)	(0.036)	(0.037)	(0.041)
Risk=1	-0.036	0.017	0.017	0.020	0.137
	(0.143)	(0.142)	(0.142)	(0.163)	(0.182)
Flood=1	0.119***	0.117***	0.117***	0.112***	0.105***
	(0.015)	(0.015)	(0.016)	(0.018)	(0.026)
Observations	31697	30891	27797	23848	14852
Students	5609	5480	4979	4331	2667
Maximum Elevation	All	100m	90m	80m	70m
Maximum Distance	1000m	1000m	1000m	1000m	1000m

The dependent variable is standard deviation in final exam score. A Risk value of one indicates an address with elevation below 60 meters. A Flood value of one indicates April 2017 exam scores. A Flood value of zero indicates April 2016 exam scores. Within-student fixed effects model. Heteroskedasticity robust standard errors are in parentheses, clustered at the student level. (*** $p<0.01$, ** $p<0.05$, * $p<0.1$.)

Table 3: The Effect of 2019 Flooding Risk on Performance

	(1) Score	(2) Score	(3) Score	(4) Score	(5) Score
Risk=1 × Flood=1	-0.070*	-0.069*	-0.071*	-0.083**	-0.113***
	(0.038)	(0.038)	(0.038)	(0.039)	(0.042)
Risk=1	0.154	0.154	0.170	0.177*	0.241*
	(0.102)	(0.102)	(0.104)	(0.107)	(0.144)
Flood=1	0.091***	0.090***	0.093***	0.105***	0.126***
	(0.014)	(0.014)	(0.015)	(0.017)	(0.023)
Observations	34477	33826	30693	26835	16063
Students	6021	5904	5403	4773	2863
Maximum Elevation	All	100m	90m	80m	70m
Maximum Distance	1000m	1000m	1000m	1000m	1000m

The dependent variable is standard deviation in final exam score. A Risk value of one indicates an address with elevation below 60 meters. A Flood value of one indicates April 2019 exam scores. A Flood value of zero indicates April 2018 exam scores. Within-student fixed effects model. Heteroskedasticity robust standard errors are in parentheses, clustered at the student level. (*** $p<0.01$, ** $p<0.05$, * $p<0.1$.)

Table 4: The Effect of Flooding Risk on Performance by Year

	(1) 2014	(2) 2015	(3) 2016	(4) 2017	(5) 2018	(6) 2019
Risk=1 × Flood=1	0.028	0.001	0.021	-0.073**	-0.033	-0.085***
	(0.031)	(0.031)	(0.030)	(0.030)	(0.030)	(0.033)
Risk=1	0.186	0.210	-0.312**	-0.036	0.031	0.162*
	(0.274)	(0.146)	(0.141)	(0.107)	(0.080)	(0.095)
Flood=1	0.048***	0.108***	0.104***	0.119***	0.086***	0.091***
	(0.014)	(0.014)	(0.013)	(0.013)	(0.012)	(0.012)
Observations	22602	23455	26607	31697	34230	34477
Students	3854	4028	4662	5608	5963	6020
Maximum Elevation	All	All	All	All	All	All
Maximum Distance	1000m	1000m	1000m	1000m	1000m	1000m

The dependent variable is standard deviation in final exam score. A Risk value of one indicates an address with elevation below 60 meters. A Flood value of zero indicates the year prior to the column title. A Flood value of one indicates the year in the column title. Within-student fixed effects model. Heteroskedasticity robust standard errors are in parentheses, clustered at the student level. (*** p<0.01, ** p<0.05, * p<0.1.)

Table 5: Student Characteristics

	(1) Female	(2) Male	(3) Domestic	(4) Foreign	(5) Below 25	(6) Mature
Risk=1 × Flood=1	-0.044	-0.072**	-0.067***	-0.035	-0.058**	-0.015
	(0.029)	(0.036)	(0.025)	(0.058)	(0.025)	(0.054)
Risk=1	-0.015	0.153*	-0.084	0.193**	0.090	-0.151
	(0.079)	(0.093)	(0.074)	(0.095)	(0.067)	(0.125)
Flood=1	0.062***	0.051***	0.067***	0.025	0.064***	-0.005
	(0.012)	(0.014)	(0.010)	(0.020)	(0.009)	(0.028)
Observations	37990	28184	50156	16018	58510	7664
Students	5060	3711	6667	2114	7452	1320
Maximum Elevation	All	All	All	All	All	All
Maximum Distance	1000m	1000m	1000m	1000m	1000m	1000m

The dependent variable is standard deviation in final exam score. A Risk value of one indicates an address with elevation below 60 meters. A Flood value of zero indicates either April 2016 or April 2018 exam scores. A Flood value of one indicates either April 2017 or April 2019 exam scores. Student sex, domestic status, and age $>$ 25 indicators from administrative data. Within-student fixed effects model. Heteroskedasticity robust standard errors are in parentheses, clustered at the student level. (*** p<0.01, ** p<0.05, * p<0.1.)

Table 6: Housing Characteristics

	(1) All Houses	(2) Apartments	(3) Before 2010	(4) New Builds	(5) Low Rent	(6) High Rent
Risk=1 × Flood=1	-0.059**	-0.036	-0.072**	-0.028	-0.079**	-0.034
	(0.029)	(0.039)	(0.029)	(0.039)	(0.032)	(0.035)
Risk=1	-0.215	0.090	-0.022	0.058	-0.248	0.024
	(0.193)	(0.093)	(0.107)	(0.146)	(0.237)	(0.080)
Flood=1	0.049***	0.060***	0.066***	0.045***	0.075***	0.048***
	(0.014)	(0.013)	(0.013)	(0.013)	(0.017)	(0.011)
Observations	30144	34322	30730	33736	19786	46388
Students	3978	4855	4192	4687	2554	6372

The dependent variable is standard deviations in final exam score. Risk is a student address with elevation at or below 60 meters. A Flooding value of one indicates exam scores from April 2017 or April 2019. A Flooding value of zero indicates April 2016 or April 2018 exam scores. An observation is a final course grade. Housing characteristics provided by the 2016 Census Dissemination Area linked to student postal code. Within-student fixed effects model. Heteroskedasticity robust standard errors are in parentheses, clustered at the student level. (*** p<0.01, ** p<0.05, * p<0.1.)

Table 7: Migration and Address Distance to Water

	(1) 2016	(2) 2017	(3) 2018	(4) 2019
Year=2016	-8.179			
	(8.062)			
Exposed=1 × Year=2016	-67.284***			
	(24.939)			
Year=2017		-9.112		
		(6.669)		
Exposed=1 × Year=2017		-66.454***		
		(20.334)		
Year=2018			-10.888*	
			(6.301)	
Exposed=1 × Year=2018			78.509***	
			(18.619)	
Year=2019				-4.780
				(8.598)
Exposed=1 × Year=2019				55.130**
				(26.482)
Observations	13970	21612	22185	15328
Students	8783	13115	13375	9598

The dependent variable is address distance to water, measured in meters. Exposed indicates students who were in a low-lying area during the 2017 floods. Each column compares the titular year against the year prior. Within-student fixed effects model. Heteroskedasticity robust standard errors are in parentheses, clustered at the student level. (*** $p<0.01$, ** $p<0.05$, * $p<0.1$.)

Table 8: Migration and Address Elevation

	(1) 2016	(2) 2017	(3) 2018	(4) 2019
Year=2016	0.047			
	(0.067)			
Exposed=1 × Year=2016	-1.182***			
	(0.207)			
Year=2017		0.005		
		(0.071)		
Exposed=1 × Year=2017		-1.562***		
		(0.216)		
Year=2018			-0.214***	
			(0.080)	
Exposed=1 × Year=2018			2.326***	
			(0.236)	
Year=2019				-0.264**
				(0.106)
Exposed=1 × Year=2019				0.779**
				(0.326)
Observations	13970	21612	22185	15328
Students	8783	13115	13375	9598

The dependent variable is address elevation, measured in meters above sea level. Exposed indicates students who were in a low-lying area during the 2017 floods. Each column compares the titular year against the year prior. Within-student fixed effects model. Heteroskedasticity robust standard errors are in parentheses, clustered at the student level. (*** p<0.01, ** p<0.05, * p<0.1.)

Table 9: Regression Discontinuity Approach

	(1) 2016	(2) 2017	(3) 2018	(4) 2019
Above $60m$ Cutoff	0.281	0.896*	-0.450	0.918*
	(0.536)	(0.525)	(0.463)	(0.472)
Threshold (m)	60	60	60	60
Obs. $<$ Threshold	1976	2024	2301	2069
Obs. $>$ Threshold	5336	7272	6979	7391

Results correspond to coefficient τ in Equation 4, the weighted local difference in the dependent variable (residual performance) around the elevation cutoff of 60 meters. Measured in percentage points. Within-student variation in percentage points is approximately 12.2 suggesting the estimated effect size in column 2 is 7.3% of a standard deviation. In column 4 the estimated effect size is 7.5% of a standard deviation. A triangular weighting kernel was used around the threshold. (*** p<0.01, ** p<0.05, * p<0.1.)

Table 10: Rain

	(1) 2016	(2) 2017 (Flood)	(3) 2018	(4) 2019 (Flood)
Total Precip (mm)	-0.009	0.000	0.002	0.002***
	(0.006)	(0.002)	(0.002)	(0.001)
Risk=1 $\times$ Total Precip (mm)	-0.012	-0.009**	-0.002	-0.005**
	(0.013)	(0.004)	(0.007)	(0.002)
Observations	11,981	13,914	13,594	13,640
Students	3,460	3,993	4,016	4,042
Proportion at Risk	0.16	0.16	0.14	0.15
Avg. Precip. (mm)	0.67	2.37	1.34	8.59
SD Precip. (mm)	1.52	4.10	3.25	9.43

The dependent variable is standard deviations in exam score. Total daily precipitation measured in millimeters. Risk is a student address with elevation at or below 60 meters. An observation is a final course grade. Each column separately estimates each year in the sample. Data from April 2016, 2017, 2018, and 2019. Proportion at risk is the proportion of exams written by those below 60 meters. Average and standard deviation of precipitation are implicitly weighted by the number of exams written by day. Within-student within-exam period fixed effects model. Heteroskedasticity robust standard errors are in parentheses, clustered at the student level. (*** p<0.01, ** p<0.05, * p<0.1.)

Table 11: Water Level

	(1) 2016	(2) 2017 (Flood)	(3) 2018	(4) 2019 (Flood)
Water Level > Elevation	-0.410***	-0.480***	-0.560***	-0.516***
	(0.082)	(0.078)	(0.098)	(0.074)
Observations	14,587	17,110	17,120	17,357
Students	3,565	4,211	4,234	4,242
Students (Flooded)	319	373	284	355
Avg. Water Level (m)	59.30	59.69	58.22	59.21
Max. Water Level (m)	59.61	59.91	58.62	60.26

The dependent variable is standard deviations in exam score. Water level and elevation measured in meters above sea level. Water level greater than elevation is an indicator for risk of flooding. An observation is a final course grade. Each column separately estimates each year in the sample. Data from April 2016, 2017, 2018, and 2019. Students (flooded) refers to the number of students who 'switch' from being unflooded to flooded status in that exam season. Average water level is implicitly weighted by the number of exams written by day. Within-student within-exam period fixed effects model. Heteroskedasticity robust standard errors are in parentheses, clustered at the student level. (*** $p<0.01$, ** $p<0.05$, * $p<0.1$.)

Table 12: Pageviews Accessing Regulator's Flood Warning Website

	(1) 2016	(2) 2017	(3) 2018	(4) 2019
Water Level (m)	1477.307***	6993.592***	315.206	22852.995***
	(274.609)	(2013.572)	(507.212)	(1726.740)
Days	15	15	20	18
Avg. Views	821	3,296	703	11,415
SD Views	342	2,548	286	15,340
Max Daily Views	1,530	7,764	1,365	45,464

The dependent variable is daily unique pageviews to the Regulator's website, which displays flooding status. Water level measured in meters above sea level. An observation is a day with a nonzero number of exams. Each column separately estimates a year in the sample. Data from April 2016, 2017, 2018, and 2019. Average number of views, the standard deviation and the max in the exam season also presented. (*** $p<0.01$, ** $p<0.05$, * $p<0.1$.)

Table 13: Users Accessing Regulator's Flood Warning Website

	(1) 2016	(2) 2017	(3) 2018	(4) 2019
Water Level (m)	233.198***	1111.009***	49.039	2873.480***
	(33.854)	(327.986)	(60.714)	(303.737)
Days	15	15	20	18
Avg. Users	123	478	117	1,576
SD Users	51	410	34	2,005
Max Daily Users	230	1,272	185	6,567

The dependent variable is daily unique users to the Regulator's website, which displays flooding status. Water level measured in meters above sea level. An observation is a day with a nonzero number of exams. Each column separately estimates a year in the sample. Data from April 2016, 2017, 2018, and 2019. Average number of users, the standard deviation and the max in the exam season also presented. (*** p<0.01, ** p<0.05, * p<0.1.)

Table 14: Flood Concern Proxy

	(1) 2016	(2) 2017 (Flood)	(3) 2018	(4) 2019 (Flood)
Pageviews	-0.035	0.013***	-0.041	0.000
	(0.024)	(0.003)	(0.026)	(0.000)
Risk=1 × Pageviews	-0.002	-0.017**	0.018	-0.003**
	(0.053)	(0.007)	(0.074)	(0.001)
Observations	11,981	13,914	13,670	13,659
Students	3,460	3,993	4,017	4,042
Proportion At Risk	0.16	0.16	0.14	0.15
Avg. Pageviews ('000s)	0.84	3.49	0.76	14.73
SD Pageviews ('000s)	0.34	2.47	0.29	15.52

The dependent variable is standard deviations in exam score. Pageviews are derived from data from the official source of flood warnings website traffic. An observation is a final course grade. Each column separately estimates each year in the sample. Data from April 2016, 2017, 2018, and 2019. Proportion at risk refers to the proportion of exams written by students with address below 60 meters. Average Pageview is implicitly weighted by the number of exams written by day. Within-student within-exam period fixed effects model. Heteroskedasticity robust standard errors are in parentheses, clustered at the student level. (*** p<0.01, ** p<0.05, * p<0.1.)

Table 15: Flood Warning Direct Effect

	(1) 2017	(2) 2019
Warning Issued=1	-0.040**	0.007
	(0.018)	(0.015)
Warning Issued=1× Risk=1	-0.021	-0.092**
	(0.041)	(0.039)
Observations	13,914	13,640
Students	3,993	4,042
Proportion At Risk	0.16	0.15

The dependent variable is standard deviations in exam score. A major flood alert was issued on April 18, 2017. On that day and the days following the binary alert variable is equal to 1 (also see Section 3.1). A flood alert (expected levels equal to those from 2017) was issued on April 18, 2019. On that day and the days following the binary alert variable is equal to 1. (also see Section 3.2). An observation is a final course grade. Each column separately estimates each year in the sample. Data from April 2017 and 2019. Proportion at risk refers to the proportion of students with address below 60 meters. Within-student within-exam period fixed effects model. Heteroskedasticity robust standard errors are in parentheses, clustered at the student level. (*** $p<0.01$, ** $p<0.05$, * $p<0.1$.)

A Appendix Figures

Figure A1: Disaster Recovery Assistance for Ontarians 2017 Eligibility Polygons

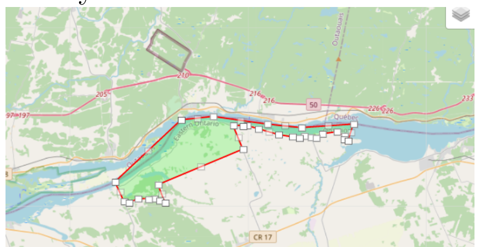

Alfred Area

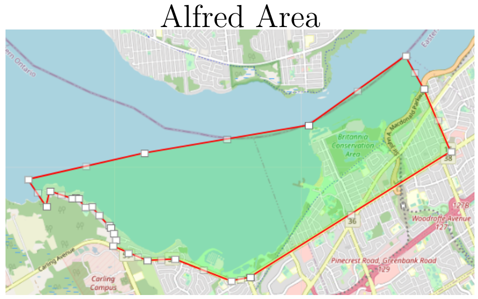

Bay Area

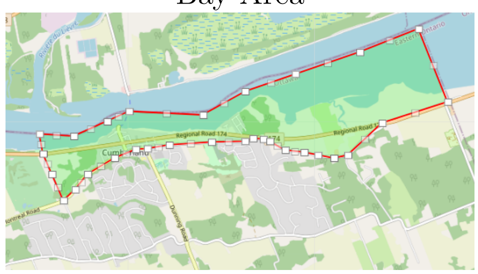

Cumberland Area

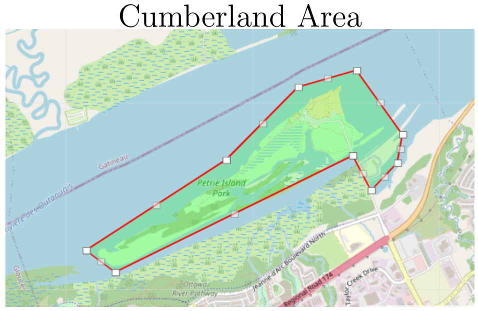

Petrie Area

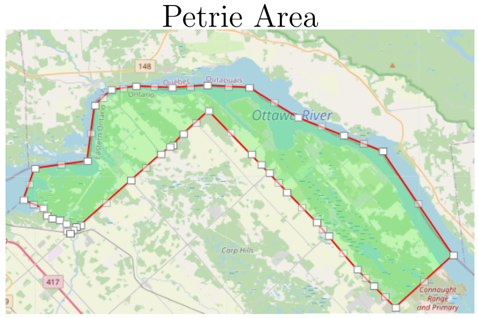

West Area

Reconstructed polygons of areas with homeowners eligible to apply for provincial government assistance under the “Disaster Recovery Assistance for Ontarians” activation for the 2017 flooding.

B Appendix Tables