Agriculture, Trade, and the Spatial Efficiency of Global Water Use

Tamma Carleton* UCSB and NBER
Levi Crews† Princeton
Ishan Nath‡ San Francisco Fed §

22 January 2024

Please click here for the latest draft.

Abstract

Over 90% of global water use occurs in agricultural production, which is subject to two pervasive distortions: (i) incomplete property rights for farmers accessing water and (ii) subsidies, taxes, and tariffs affecting agricultural output. This paper combines a rich collection of global geospatial data with a dynamic spatial equilibrium model to quantify the impact of agricultural and trade policies on regional water scarcity and welfare. In the data, we show that water-intensive crops concentrate highly in water-abundant locations, implying a strong role for comparative advantage in governing global water use, though a small number of regions with very water-intensive production are losing water rapidly over time. In the model, we capture production, consumption, and trade in agriculture across many countries and crops, as well as the dynamic evolution of local water stocks as farmers extract from the common pool resource. We calibrate the model to match observed global patterns of agricultural production and hydrologic trends and use it to conduct counterfactual simulations of alternative policy regimes. We find that eliminating international trade in agriculture would dramatically increase water depletion across most of the world, and especially so in drier food-importing regions, resulting in large reductions in welfare over time. In contrast, other observed and hypothetical agricultural policy liberalizations have mixed effects on depletion that vary greatly across locations, suggesting nuance in implications for policy.

JEL codes: F13, F18, H23, Q15, Q17, Q25, Q27

---

* tcarleton@ucsb.edu
† lcrews@princeton.edu
‡ inath33@gmail.com

§ We thank Zoë Arnaut-Hull, Cheikh Fall, Yang Gao, Ethan Goode, Christian Hilgemann, Walker Lewis, and Samuel Nitkin for excellent research assistance. We also thank seminar and conference participants at UBC, USC, UC Berkeley, the World Bank, Yale, the Coase Project Conference, the Fed Junior Applied Seminar, the Latin American and Caribbean Economics Association Annual Meeting, and the Virtual Workshop on Trade, Spatial Economics, and the Environment for helpful comments. Any views expressed in this paper are those of the authors and do not necessarily reflect those of the Federal Reserve System or its staff.

1 Introduction

Over 90% of global water use by humans occurs in agricultural production, which is critically dependent on rainfall and local stocks of groundwater and surface water (Mekonnen and Hoekstra, 2011). While water itself is generally prohibitively costly to transport over long distances, nearly 25% of all water consumption is embedded in internationally traded agricultural products. Across the world, both input markets for water and output markets for agriculture are subject to pervasive distortions. Most farmers extract water as an open access resource without defined property rights (Libecap, 2008), and agricultural markets typically exist amidst a broad array of subsidies, taxes, tariffs, and trade restrictions (Anderson, Rausser and Swinnen, 2013). When input market failures prevent the cost of water from reflecting its scarcity, trade liberalization can exacerbate the impact of this distortion and have adverse long-run effects on resource depletion and welfare (Chichilnisky, 1994).

A few prominent examples suggest that global trade in agriculture could be contributing to severe regional water depletion. In California's Central Valley, groundwater stocks have declined precipitously in recent decades in places that specialize in producing and exporting highly water-intensive agricultural goods. For instance, California produces approximately 80% of global almonds, which are the world's second most water-intensive crop per acre of production. Similarly, India's northern agricultural regions have been losing water faster than almost any other arable land on earth (Rodell et al., 2018) while cultivating water-intensive crops like rice, of which India is the world's leading exporter. Sekhri (2022) shows that in this context, policies encouraging trade led directly to substantial groundwater depletion.

Beyond such examples for a few specific locations, no consensus exists on the systematic global effects of international trade on long-run water resources and agricultural production. In the scientific literature, prominent work by Dalin et al. (2017) finds that the vast majority of traded agricultural goods are produced in regions depleting their groundwater resources. In contrast, classic work by the geographer Tony Allan (Allan, 1998, 2011) argues that "global trade enables local water security" by allowing production in water-abundant locations to support consumption in drier regions. Notably, despite the critical role of international agricultural markets in this question, existing research has taken place almost entirely outside of economics. Achieving an understanding of the effects of global trade and agricultural policies on water resources has been held back both by the historical lack of granular systematic data on water resources, and by the absence of analysis using frameworks that incorporate equilibrium behavior, long-run dynamics, and welfare.

This paper combines a rich collection of global geospatial data with a quantitative dynamic spatial equilibrium model to analyze the effects of global trade and agricultural policies on regional water scarcity and long-run welfare. We start by compiling a novel combination of globally comprehensive datasets that convey a wide range of information about water resources, agricultural activity, and agricultural policy. For water, two datasets recently assembled in the natural science literature use a combination of direct observations from well monitors, hydrological modeling, and satellite observations to provide a high-resolution representation of global groundwater table depths, surface water presence, and trends in total water resource stocks in recent decades. We combine these data with gridded information on rainfall, aridity, and soil characteristics relevant to modeling the evolution of local water accessibility over time. For agriculture, we collect high-resolution global data on harvested acreage by crop, pastureland, the share of land equipped for irrigation, and potential yields of each crop in each location based on local soil and climate characteristics, along with crop-level data on water-intensity of production and country-crop-level data on production, prices, trade, and a summary measure of the net effective tax or subsidy from a broad range of policy interventions.

We use the data and corresponding scientific literature to establish a set of five facts about water, agriculture, and trade that frame the analysis. We start by showing first that water resources exhibit tremendous spatial heterogeneity across the globe. Some regions have ample access to groundwater, surface water, and rainfall, while others have both difficult-to-access and declining volumes of water resources, implying a critical potential role for spatial reallocation in mediating the global costs of water scarcity. Second, we emphasize the literature showing that agriculture dominates human water use, which motivates the sectoral focus in this paper. Third, we show that the vast majority—over 93%—of the world's agricultural production occurs in locations where farmers use water as an open-access resource with no formal or tradable property rights, underscoring the importance of understanding the impact of output market and trade policies in the presence of this ubiquitous input market failure. Fourth, we present direct empirical evidence from previous work on the substantial effects of agricultural and trade policies on the evolution of water resources.

The fifth and most important fact we present shows that water-intensive agricultural activity is, on average, highly concentrated in locations with abundant water resources. For example, the average water-intensity of agricultural activity on arable land in regions in the top global quintiles of groundwater availability and rainfall is two and five times higher, respectively, than in regions in the bottom quintiles. The relatively water-intensive use of water-abundant regions is driven both by cropping a larger proportion of the acreage, and by choosing more water-intensive crops conditional on planting in a given location. Consistent with previous work in particular regions, the data do also show a small number of regions losing water rapidly while engaging in highly water-intensive production, such as in parts of California and the state of Uttarakhand in north India. What is perhaps surprising, and to our knowledge has not been previously established, is that these regions constitute an extreme exception to the systematic global pattern. Overall, the facts suggest that drier regions of the world preserve their limited water resources through a combination of producing low water-intensity crops locally and importing more water-intensive forms of agricultural production.

To fully characterize the role of trade in governing global water resources and explore policy counterfactuals, we next build a dynamic spatial equilibrium model of water resources and agricultural production, consumption, and trade across many crops and countries. In the model, farmers on each parcel of arable land on earth choose whether to plant crops or to work in the non-agricultural sector. They sell their output in domestic and foreign markets subject to crop-specific subsidies, taxes, and bilateral trade costs. The productivity of growing each crop depends on local soil quality, climate, and the cost of extracting the crop-specific requirement of water. Extraction costs depend on local water scarcity, and the stock of water in each local aquifer evolves dynamically depending on natural recharge and agricultural land use, thus embedding the spatial and temporal externalities caused by open access to the resource. If aquifers draw down over time, crops become more difficult to produce, agricultural productivity declines, food prices rise, and welfare suffers.

We use the extensive global dataset to calibrate the model for a broad range of countries that account for 99% of the world's agricultural workers. Farmers in the model operate across approximately two million granular "fields" of land that represent heterogeneous local soil and climate endowments, split across over 200 global "aquifers" that reside within and across countries. We do not impose a steady-state in the baseline calibration, such that the model reproduces both levels and current trends in regional water resources. The model simulations include 22 crops that range from globally traded staples such as wheat, rice, maize, soybeans, and potatoes, to specialty water-intensive cash crops such as coffee, oil palm, and bananas, to regional crops critical in many drier low-income regions, such as cassava, sorghum, millet, barley, and chickpeas. We calibrate some model parameters from the literature, such as using scientific estimates of crop-specific water intensity of production, and others to match observed data on land use, agricultural production, and water resource levels and trends.

We use the calibrated model to run three sets of counterfactuals. First, we evaluate the fundamental question about the effects of global agricultural trade on water resource depletion by considering what the world would look like in its absence—i.e., simulating a scenario that imposes autarky on all countries and crops starting in our baseline year. Second, we quantify the effects of historical realized policy reforms by simulating the impacts of the global agricultural and trade policy liberalizations prompted by the Uruguay Round of World Trade Organization negotiations that occurred from 1986-1994. We consider both the effects of the reforms that were implemented in the decades following the signing of the bill by 123 nations in 1994, as well as the effects of hypothetical further changes that eliminate remaining distortionary agricultural policies across the world. Finally, in future drafts of the paper we will consider the domestic and global water resource and welfare implications of stylized versions of two observed unilateral country level policies: an export ban on rice in India, a critical food exporter that regularly uses export bans to insulate domestic staple crop prices from global supply shocks, and an import substitution strategy characterized by import restrictions and domestic agricultural subsidies in food importing countries such as Egypt and Turkey.

The paper has three key findings. First, we find that existing global trade dramatically reduces aggregate global land and water use, preventing substantial dynamic welfare losses from water depletion. Global cropped acreage nearly doubles in autarky in order for all countries to meet domestic demand in the absence of the efficiency gains from trade, which channels agricultural activity toward its most productive global locations. The worldwide share of arable land that is cropped goes from about one-third in the baseline, which closely matches the data, to nearly two-thirds in autarky. To our knowledge, these represent the first quantitative estimates of the effects of global trade on agricultural land use, which could have implications for other environmental issues beyond the scope of this paper.

Autarky also raises global agricultural water use by about 60% in the initial period. This sharp increase in global extraction depletes global water resources substantially over time. Average water table depth across the world's aquifers is stable over time in the baseline simulation, which matches the data, but falls by about 27% in the first 30 years of autarky as extraction persistently outstrips recharge. The decline in global freshwater availability in autarky causes considerable welfare losses. As water tables fall, crop yields decline, food prices rise, and production in the outside sector falls as agriculture commands a greater share of productive resources. We calculate that, when normalizing welfare in the initial period to its baseline value to net out the static gains from trade, global welfare in autarky falls by about an additional 2.2% after 30 years due to water resource depletion. Thus, we estimate that the previously unmeasured dynamic water resource gains from trade are of the same order of magnitude as traditional measures of the global static gains from trade (Arkolakis, Costinot and Rodríguez-Clare, 2012).[1]

The second key finding is that water-scarce regions would suffer the worst long-run water depletion and welfare losses in the absence of global agricultural trade. The 30-year decline in water tables in autarky is over three times larger for aquifers initially at the 90th percentile of depth than for the global average. These water-scarce locations primarily consist of food-importing countries, which require the largest increases in cropped acreage and water use to sustain domestic consumption in the absence of trade. For these regions with little rainfall and low water tables, the existence of trade preserves their long-run access to water and prevents substantial welfare losses—as much as 10-15% of initial consumption—from water depletion over time. Consistent with previous empirical work (Sekhri, 2022) and circumstantial evidence, the model simulations do show that imposing autarky *reverses* the trend of declining water resources in a small number of rapidly depleting food exporting regions, including California's Central Valley and Northern India,

---

[1] Note that the welfare calculations in this draft of the paper remain preliminary, and are subject to change.

which stand out as exceptions in the global distribution of water changes under autarky. In India, welfare is declining over time in the *baseline* as it depletes its water resources, but is dynamically stable in autarky when it can no longer deplete its water stocks to export water-intensive products. For the overwhelming majority of the world, however, we find that trade in agriculture allocates water-intensive production to water-abundant locations, dramatically increasing the spatial efficiency of production and preserving long-run water availability.

It is worth noting that the strength with which specialization in water-intensive agriculture follows the resource's abundance, and the degree to which trade prevents long-run depletion, is perhaps surprising given not only the prominent examples of trade-induced water declines in California and India, but also the pervasive market failures in procuring water as an input. If there are no property rights, no markets, and no market prices for the vast majority of producers using this input, what is the mechanism by which its relative abundance maps into specialization? We show that the economic explanation for this finding is that water scarcity directly governs its extraction cost. In calibrating the model, we infer the heterogeneous costs of extraction in each aquifer from the revealed preference of farmers in the data choosing how water-intensively to use their land. We show that, on average, input costs for water are substantially higher in locations with lower water tables and less rainfall. Thus, despite the lack of formal markets or tradable property rights, water's *effective* input price *does*, at least partially, reflect its scarcity. In this setting, the natural environment broadly stands in for functioning input markets by raising the marginal cost of procurement where water is scarce.

The third key finding is that while existing trade generally productively reallocates water use across space, specific agricultural policy liberalizations can have the opposite effect and exacerbate depletion. In the counterfactuals we consider, we find that the Uruguay Round of WTO negotiations, the largest historic liberalization of agricultural markets to date, led global agriculture to, on average, reallocate from water-abundant to water-scarce countries, with corresponding implications for long-run water availability and welfare. When agriculture became broadly subject to WTO rules, many wealthier, water-abundant countries substantially reduced domestic support provided to the agricultural sector, while many lower-income countries removed disincentives to agriculture that were related to import substitution policies designed to promote industrialization (Anderson, Rausser and Swinnen, 2013).

Our model simulations show that these policy reforms reduced water use in many water-abundant regions, such as in Western Europe, and increased it in many scarce locations, such as in parts of Sub-Saharan Africa. In the long-run, these reforms appear to have slightly exacerbated depletion on average across the world's aquifers, and especially so in a large number of relatively dry countries. In contrast, we show that if remaining domestic agricultural market distortions were removed, nearly all of which subsidize agricultural production, depletion risks would decrease slightly both globally and in most dry locations. In general, the results suggest that the long-run water implications of agricultural and trade policy reforms can vary substantially across cases, and that reducing output market distortions need not be beneficial in the long-run in this setting.

This paper builds on related work across several literatures within and beyond economics. In the literature on the economics of water, most existing papers focus on the efficacy of input market reforms and related policies that take place in a single location, such as California, India, or southeastern Australia (Ayres, Meng and Plantinga, 2021; Bruno and Jessoe, 2021; Ryan and Sudarshan, 2021; Rafey, 2023). A small number of empirical papers consider the role of international trade in agriculture, including Carleton (2021), which estimates the impact of agricultural and trade policies on trends in water resources, and Sekhri (2022), which investigates trade promotion and groundwater depletion in northern India. The paper with the most relevant empirical work to the facts presented here is Debaere (2014), which uses country-level data to show that water-intensive exports correlate with water-abundance, though less so than is observed for other inputs.

To our knowledge, this paper is the first quantitative spatial analysis of the allocation of global water resources. This work builds on a growing recent economics literature at the intersection of trade, spatial, and the environment, which is summarized in a review article by Copeland, Shapiro and Taylor (2022). The model in this paper builds most closely on Costinot, Donaldson and Smith (2016), though we add water resources and dynamics. Most papers in the spatial environmental literature focus on climate change and air pollution, but a small number consider natural resources in the context of forests. Hsiao (2021) investigates the effects of import tariffs on deforestation for palm oil production in Indonesia and Malaysia, and Domínguez-Iino (2021) considers the land use and deforestation implications of market power in agricultural supply chains in South America. The most similar work to this paper is a new paper by Farrokhi, Kang, Pellegrina and Sotelo (2023), which uses a global dynamic spatial equilibrium model to study the effects of agricultural trade on deforestation. The context of forests differs from water in a number of critical ways, including most notably that the CO_2 externality is global rather than local, and the marginal cost of extraction varies little across space.

As mentioned above, this paper also relates to an older theoretical literature on trade and natural resources, and a number of papers on "virtual water trade" published in scientific journals. In economics, Chichilnisky (1994) most prominently uses a simple two country model to make the qualitative point that the welfare effects from trade liberalization are ambiguous in a setting with distorted input markets for natural resources. In the natural sciences, a number of papers consider the implications of "virtual water trade" for a wide range of topics such as groundwater depletion (Dalin et al., 2017), inequality across countries (Carr et al., 2015), and climate change (Konar et al., 2013). Since these papers do not contain economic models of supply, demand, and trade, however, they do not quantify counterfactual policies, equilibrium reallocation, long-run dynamics, or welfare.

The paper proceeds as follows. Section 2 describes the wide array of geospatial data compiled in this paper, and Section 3 uses it to establish a set of stylized facts that frame the analysis. Section 4 lays out the model, and Section 5 shows how we calibrate it to match the data. Section 6 shows results from the policy counterfactuals, and Section 7 concludes.

2 Data

We start by compiling a wide array of geospatial datasets to assemble what constitutes, to our knowledge, the largest collection of global data on water and agriculture yet to be used in economics.

2.1 Spatial Data on Global Water Resources

Our analysis of the spatial allocation of water resources draws primarily on two scientific datasets novel to the economics literature that provide information on levels and trends of water availability throughout the world.

Groundwater Table Depth

We collect globally comprehensive data on a cross-sectional measure of water table depth from Fan, Li and Miguez-Macho (2013).[2] This scientific paper begins by compiling water table depth observations published in government or scientific sources from over 1.6 million wells located across all six populated continents. The paper proceeds to create a continuous spatial dataset by interpolating between well observations using a hydrological model calibrated to detailed spatial data on climate, geology, elevation, and soil characteristics. Together, the empirical observations and model simulations are used to produce global estimates of the depth of the water table in meters from the surface at a 30 arc-second (approximately 1km) resolution.

This data provides estimates of the water available to farmers from both groundwater and surface water. For farmers who irrigate their crops using groundwater, well depth plays a critical role in their costs of extraction since it is costlier to pump water from further underground.[3] Other farmers irrigate crops from surface water, which is typically lower cost where available. Approximately 37% of global irrigated land relies on groundwater, and 63% on surface water.[4] The data from the hydrological model simulations in Fan, Li and Miguez-Macho (2013) provides information on the presence of surface water by marking these areas with water table depth readings of less than zero. The paper estimates that 15% of global land area is

---

[2] Note that we use the updated version of the dataset accessed here, which corrects for some known errors in the 2013 version.

[3] See Hendricks and Peterson (2012) for analysis of how extraction costs vary with well depth.

[4] This statistic comes from our calculations using data from the Monthly Irrigated and Rainfed Crop Areas (MIRCA2000) dataset produced by Portmann et al. (2010).

covered by lakes, rivers, and inundated wetlands (this proportion excludes oceans, seas, and other large water bodies such as the Great Lakes).

The data on both water table depth and surface water availability plays a critical role in the calibration of the costs of water extraction that follow in this paper. The drawback of this dataset, however, is that it contains only cross-sectional information on water availability, limiting its ability to inform about water resource dynamics that exist in the world and are represented in the model.[5]

Trends in Total Water Availability

The second key water dataset in this paper contains globally gridded information on trends in the total volume of local water resources over the period from 2003 through 2016.[6] The data is collected by the Gravity Recovery and Climate Experiment (GRACE) partnership between the U.S. National Aeronautics and Space Administration (NASA) and the German Deutsche Forschungsanstalt für Luft und Raumfahrt (DLR). The scientific procedure underlying this data exploits satellite measurements of the gravitational force exerted by each location on Earth to infer its change in mass, and consequently total water storage, over time. Tapley et al. (2004) show that the time-series changes in regional mass uncovered by GRACE primarily consist of changes in local water content.[7]

The GRACE data provide month-by-month changes in total water storage (TWS) measured in centimeters of equivalent water height[8] for 41,168 equal-area parcels of the Earth, which measure 1 1 at the equator. Because the data infers the change in all water contained in a given region, the trends represent the combined evolution of groundwater, surface water, soil moisture, snowpack, and ice over time (though the latter two components are less relevant on the arable land that is the focus of this paper). This data has been used widely in the scientific literature to study topics ranging from ice sheets to ocean currents to groundwater storage on land, but has been almost completely absent from economics.[9]

The primary drawback of the GRACE data is that the measurement of trends inferred by the satellite measurements provides no information about the level of groundwater or surface water in any location at any point. Thus, we couple the time-series data on trends from GRACE with the cross-sectional data on levels of groundwater and surface water from Fan, Li and Miguez-Macho (2013) to provide

---

[5] The exact years of water table depth observations in Fan, Li and Miguez-Macho (2013) vary across regions, but end in 2009.

[6] This time period corresponds to the original GRACE satellite mission. A follow-on mission continues to extend this time frame, but we limit our analysis to the original mission for consistency reasons and because these years overlap with other agricultural data used.

[7] See Carleton (2021) for a more complete description of the GRACE data.

[8] Equivalent water height is defined as the depth of water that would be present were it to be spread evenly across the entire surface of a grid cell.

[9] The two exceptions to this that we are aware of are recent work by Taylor (2022) and a related preceding paper by Carleton (2021).

a more complete picture of the recent status and dynamic evolution of water availability throughout the world.

Other Hydrological Data

We supplement the two primary datasets on water availability described above with a range of other relevant global spatial hydrological datasets. We use data on cumulative precipitation at 0.25 × 0.25 resolution from the Global Meteorological Forcing Dataset (GMFD) version 3 from Princeton University (Sheffield, Goteti and Wood, 2006). We also collect satellite data on the presence of surface water at a 30 meter resolution from Pekel et al. (2016) as a complement to the measure available in the Fan et al. (2013) data.

We collect several datasets relevant to calibrating the law of motion for local water stocks in the model. These include global spatial data at the 30 arc-second resolution on aridity, a measure of evapotranspiration potential that governs the proportion of rainfall that contributes to recharging local water stocks. This data is produced by Trabucco and Zomer (2019) and made available by the CGIAR Consortium for Spatial Information. In addition, data on soil type from Hengl et al. (2017) and specific yield by soil type from Loheide et al. (2005) contribute to mapping changes in the volume of water in a given location to changes in water table depth in the model simulations.

2.2 Global Agricultural Datasets

We collect global spatial data on agricultural land use and potential crop productivity, crop-level data on water use, and country-level data on agricultural production, consumption, trade flows, and policy interventions. The details are as follows.

Agricultural Land Use

We use global spatial data on agricultural land use at a 10km × 10km resolution compiled by Monfreda, Ramankutty and Foley (2008). These data use a combination of remote sensing and census records from national and subnational government entities to estimate the fraction of land area in each grid cell allocated to planting each of 175 crops, and to pasture land, in the year 2000.[10] In this paper, we use the union of cropped area and pasture land as the definition of arable land with the potential for cultivation.

We supplement the data on crop choice and land use with additional spatial data that measures the proportion of agricultural land equipped for irrigation. This dataset is called the Global Map of Irrigation Areas, and is compiled by the

---

[10] While more recent products exist for aggregate cropped area, such as Potapov et al. (2022), updated estimates of crop- and pasture-specific areas are not available.

United Nations Food and Agriculture Organization (FAO) and Rheinische Friedrich-Wilhelms University. The data comes at the 5 arc-minute resolution and contains estimates for the year 2005.

Potential Yields

The second global spatial agricultural dataset comes from the FAO's Global Agro-Ecological Zones (GAEZ). This data contains potential yields for 38 crops at 5 arc-minute resolution, estimated using an agronomic model that incorporates detailed local information on soils, geography, and climate. The model provides crop-specific estimates of the maximum yield attainable under a range of possible assumptions about farmer inputs and climate conditions.

Critically, this data provides estimates of potential yields for all crops in all locations, including those that have not been historically observed in a given place. This allows for a rich representation of regional comparative advantage in agriculture and for counterfactual model simulations in which crop choice can shift meaningfully across location. Costinot et al. (2016) pioneered the use of this data in quantitative trade models. In our implementation, we follow their work in using the high-input yield estimates, but restrict attention to the historical climate scenario that uses average weather from 1961-1990. We take the weighted average of "rain-fed" and "irrigated" potential yields, using the data on the area equipped for irrigation in each location to assign weights between the two.

Crop Specific Water-Intensity of Production

We use data from Mekonnen and Hoekstra (2011) to estimate the average global water-intensity of each crop. These estimates are denominated in units of m^3 of water used as an input per ton of output for each crop. We combine them with data on the average global yield of each crop to construct a dataset of crop-specific water intensities denominated in units of m^3 per hectare. Note the caveat that Mekonnen and Hoekstra (2011) provide only an average global estimate of crop water use without any heterogeneity, so a data limitation of the analysis is that we cannot account for changes in water use arising from different techniques of growing a given crop.

Production and Trade

We collect data from FAOSTAT on crop-specific production in metric tons along with farm-gate prices measured in USD per ton. The data is available for over 200 countries from 1961-2020, though in our implementation in the model calibration we use a cross-section of the data for a subset of countries in 2009. We also use data on bilateral trade flows by crop for the same year from the UN Comtrade database.

Agricultural Policy

Government interventions in agricultural markets play a critical role in this paper. Analyzing these policies in an international context, however, is complicated by the wide array of relevant policy tools with interacting and overlapping effects, including output taxes and subsidies, input subsidies, import tariffs, quotas, sanctions, and regulations. Furthermore, these policies are implemented within institutional contexts that differ substantially across countries, such that their definition and interpretation may not be consistent across locations.

We overcome these challenges by using data from the World Bank's "Distortions to Agricultural Incentives" (DAI) project, which constructs an internationally comparable measure of agricultural policy interventions that is both comprehensive and parsimonious. Included in these data is a single summary statistic, known as the "Nominal Rate of Assistance" (NRA), measured for 80 products in 82 countries.

The NRA captures the equivalent product-specific net subsidy or tax that results from the combined effect of the full range of policies that include direct taxes and subsidies, tariff and non-tariff barriers to trade, and government manipulation of foreign exchange markets (Anderson et al., 2008). The NRA measure - positive for net subsidies and negative for net taxes - can be interpreted as the percentage difference between domestic farm-gate prices and international prices for the same product, excluding transportation and distribution costs. Critically, the measure does not include any water-specific policy interventions, such as subsidies for agricultural energy use or irrigation.

The DAI data has been used previously to study topics ranging from political economy (Anderson et al., 2013) to agricultural productivity (Adamopoulos and Restuccia, 2014). In this paper, we use this data to investigate the spatial correlation between agricultural policy and water resources in Section 3, and to calibrate the output market distortions in the model in Section 4.

3 Stylized Facts

3.1 Five Facts About Water and Agriculture

We begin the analysis by bringing together the data presented in Section 2 with a synthesis of the relevant scientific and institutional context to establish five key facts about water and agriculture that frame the quantitative model analysis to follow.

Fact 1: Water resources exhibit tremendous spatial heterogeneity

We start by summarizing the global data on water resources. Figure 1 shows the global distribution of average annual rainfall, groundwater table depth, and trends in total water storage across all the arable land on earth. The maps show enormous variation in both the levels and trends of water available in each location.

Annual average rainfall varies by almost two orders of magnitude across the world's land that is suitable for crop production. The median parcel of global arable land receives 7,850 m^3 per hectare of rainfall per year, and the 1st to 99th percentile ranges from 500 to 33,550 m^3/ha. For context, the most water-intensive staple crop, rice, requires about 8,790 m^3/ha of water per year on average, and the most water-intensive specialty crop, almonds, use about 22,200 m^3/ha.[11] Less than 45% of the world's arable land receives enough rainfall to supply full coverage of rice acreage without drawing down their stock of water in steady-state, and less than 9% receives enough to do so with almonds.[12] In contrast, the driest 36% of global arable land receives less rainfall than is required even to cover all acreage with barley or millet, the least water-intensive staple crops, which use 5,080 and 4,300 m^3/ha, respectively. To be sure, these proportions do not represent a binding limit on the share of regions that can grow each of these crops since farmers can plant on a subset of acreage or irrigate their crops from existing stocks of surface or groundwater, but average water consumption by crop provides a useful benchmark for contextualizing global rainfall totals.

Existing stocks of groundwater show similarly sharp heterogeneity. In addition to the 15% of the world's surface area that Fan, Li and Miguez-Macho (2013) show are covered by local bodies of surface water (e.g. small lakes, rivers, and wetlands) that can be used for irrigation, 19% of arable land area has easily accessible groundwater within one meter of the earth's surface. On the other hand, a large share of potentially cultivable land lies above groundwater that is difficult to prohibitive to access for irrigation. Just over 62% of the world's arable land exists in locations with groundwater deeper than eight meters from the earth's surface, a cutoff below which irrigation extraction costs become discontinuously more expensive (Sekhri, 2014). Note that while the data on groundwater depth identifies only the top most part of the water table, this information is generally sufficient to represent both the costs of extracting water for farmers and the usable stock of groundwater available since the bottom of the water table predominantly lies far below the level feasible for extraction.

The dynamic measure of trends in total water storage also varies widely across the world's arable land, previewing future availability if current patterns continue. In total, over the period from 2003 through 2016 just under 50% of the world's arable land was losing water (recall that the satellite-based measure accounts for all water in

---

[11] The data on water use by crop comes from Mekonnen and Hoekstra (2011), as explained further in Fact 3.

[12] These and related statements are highly conservative estimates, as they implicitly assume perfectly efficient agricultural water consumption. When irrigation and/or rainfall water is applied to an agricultural area, some fraction of that water is used productively by the crop, but a substantial amount of it is lost to runoff. No comprehensive estimates of runoff exist, so the statistics we report in this section conservatively assume perfectly efficient water use (and therefore no runoff). Accounting for runoff would imply that fewer regions of the globe have sufficient rainfall to supply the most water-intensive crops.

a region, including both groundwater and surface water), while the other half was gaining water. Some regions show sharp differences across subnational locations, whereas others have more consistent broader regional patterns. For instance, the southwestern U.S., western Australia, and northern India are all losing water, while the midwestern U.S., eastern Australia, and central India are all gaining. On the other hand, western Europe is largely homogeneously gaining water, while the vast majority of the Middle East and Argentina are losing water.

It is worth noting that the magnitudes of trends in total water storage are rather modest across the world. The 1st to 99th percentile of arable land ranges from losing 218 m^3/ha to gaining 197 m^3/ha. For context, median global rainfall on arable land is about 7,850 m^3/ha. Thus, the data implies that, on net, the time trends of water available in each location are small relative to the gross flows of rainfall and water consumption.

In general, heterogeneity in input abundance across locations governs the strength of comparative advantage and the potential gains from trade. Given the considerable variation in the static and dynamic measures of water availability across the world, the data suggests a strong role for the spatial allocation of production in water-intensive tradable goods in maximizing the present and future value of the world's water resources.

Fact 2: Agriculture dominates global human water consumption

Agricultural production accounts for approximately 70% of global water withdrawals by humans, including withdrawals from surface, ground, and soil for use in industrial production, energy generation, and manufacturing (Dubois et al., 2011). However, agriculture is responsible for an even larger share of the water actually *consumed* by human activity, because agricultural crops evapotranspire a large share of applied or precipitated water, such that it is lost to the local environment. In contrast, other water-extracting activities, like power plant cooling, withdraw significant quantities of water, but immediately return them to streams, aquifers, and soils. Reflecting this, leading scientific evidence indicates that 92% of global *consumptive* water use by humans is dedicated to agricultural production (Hoekstra and Mekonnen, 2012).

Hoekstra and Mekonnen (2012) traces water used in agricultural production throughout its supply chain. Specifically, the paper quantifies agricultural water consumption throughout the world at the 5 arc-minute resolution using a model that combines hydrological and agronomic mechanisms with detailed spatial data on climate and soil conditions, crop planting and harvesting dates, irrigation techniques, information on farmer inputs such as nitrogen, and the Monfreda, Ramankutty and Foley (2008) dataset on land use. Animal products account for approximately one-third of global water use when accounting for crops grown as feedstocks, though only 0.5% of that total comes from the water directly consumed by livestock. Industrial production accounts for only 4.4% of global water use, and domestic water supply for direct commercial and residential use for only 3.6%. The quantitative findings from this paper are corroborated in a recent review article by d'Odorico et al. (2019), which states that multiple scientific papers have concluded that approximately 90% of human consumptive use of water arises from agricultural production.

Hoekstra and Mekonnen (2012) and Carr et al. (2013) both also estimate that agricultural goods traded across country borders account for 20-25% of global water consumption. Relatedly, Dalin et al. (2017) calculate that 11% of groundwater extracted in excess of natural recharge is embedded in traded agricultural products. Thus, the water used to produce the consumption of the average person in a given country can differ substantially from the per capita domestic water use in production. For instance, Hoekstra and Mekonnen (2012) estimate that over 70% of implied water consumption is embedded in imported products for consumers in several Middle Eastern countries, including Lebanon, Yemen, Israel, Jordan, and Kuwait. Overall, the findings from the scientific literature suggest that analyzing the allocation of global water supplies requires a strong focus on agricultural markets and a critical role for international trade.

Fact 3: Local markets with tradable water rights rarely exist

Water is a classic example of a common pool resource in which open access to extraction creates externalities. When property rights are not clearly defined and farmers draw down the local stock of water, they raise the cost of extraction for other farmers with access to the resource both in the present and future. Thus, the social costs of using water exceed the private costs. Furthermore, even when property rights are defined, they may be allocated on the basis of historical presence or other institutional arrangements without a clear market-based mechanism for allocating the resource to its highest value uses.

Research has shown that the implementation of environmental markets that allow for trade in water property rights have large benefits. For instance, Rafey (2023) quantifies the benefits of trading in the world's largest water market in the southern Murray-Darling Basin in Australia. Similarly, Ayres, Meng and Plantinga (2021) and Bruno and Jessoe (2021) study water markets in California's Mojave Desert and Coachella Valley, respectively. Each of these papers finds substantial gains in total surplus from allowing users to trade water rights. The presence of well-defined property rights and input markets allows water to flow to those uses with the highest value, reducing the total extraction necessary to achieve a given level of output.

Despite the theoretically and empirically documented benefits of tradable property rights for water, the existence of such markets remains exceedingly rare. We conduct an extensive review of the global status of water property rights, and find that at least 94% of the world's agricultural output occurs in regions with no formal mechanisms for farmers to trade water. In addition to the water markets in Australia and California described above, we find that formally established tradable property rights for water also exist in Chile, Mexico, Spain, South

Africa, Oman, and a number of other states in the western U.S.[13] Taken together, these regions account for about 6% of the world's agricultural production value under a conservative estimate that assumes these water markets cover all agricultural activity in a jurisdiction (which is not true in most cases).

We find that some additional water trading occurs outside the small number of established formal markets, though without clear mechanisms to resolve the common pool resource challenge. For example, informal water trade has been documented in parts of India, Pakistan, and China (Saleth, 2004; Zhang et al., 2008; Easter et al., 2018).[14] While the potential existence of trade might improve allocative efficiency in these regions to some degree, however, there is no mechanism in the informal markets for restricting the total water allocation to a level at which the price reflects the resource's scarcity value. On the other hand, some locations, such as India, have implemented rationing schemes that reduce overextraction without achieving gains in allocative efficiency (Ryan and Sudarshan, 2021).

Even in those limited locations that have established formal property rights intended to optimize the use of water, implementation challenges remain prevalent. In many regions, transaction costs remain high enough that little trade occurs in practice. While the volume of water traded during droughts accounts for over 20% of usage in Australia and Chile, market trading volumes constitute only 5% of extraction in Spain and less than 1% in South Africa (Grafton et al., 2020).[15] Debaere et al. (2014) and Easter, Rosegrant and Dinar (2018) document the wide array of institutional challenges that face even the best efforts to establish successful local water markets, which have been implemented with varying degrees of success across the locations in which they exist.

Overall, our review of the global institutional context suggests that the market failures affecting input markets for water remain broadly intractable in most places. Research suggests that a few local water markets have yielded large benefits, particularly those in southeastern Australia and the western U.S., but the widespread implementation of such successful programs remains elusive. If tradable property rights schemes in which the price of water reflects its social value could be successfully scaled up across the world, analyzing challenges associated with water scarcity at a local scale would be sufficient. The first best policy in the model presented in Section 4 would be to eliminate both input side market failures and all distortions

---

[13] Other U.S. states with formal water markets for farmers include Arizona, Colorado, Idaho, Montana, Nevada, New Mexico, Oregon, Texas, Utah, Washington, and Wyoming (Griffin and Characklis, 2011; Phillips et al., 2020; Schwabe et al., 2020). For more information on the other water markets, see Easter and Huang (2014) and Grafton et al. (2020) for broad global overviews, Donoso (2012) for more details on Chile, Hearne and Trava (1997) and Kloezen (1998) on Mexico, Palomo-Hierro et al. (2015) on Spain, and Young (2012) on Australia.

[14] China has also developed a national China Water Exchange in which regional public authorities can trade water amongst themselves, and smaller pilot programs in six provinces that allow for farmer-to-farmer water exchange (Wang and Yang, 2018).

[15] Based on similar measurements, Rafey (2023) estimates that traded water rights account for a maximum of 1% of the world's freshwater withdrawals.

affecting output markets. As long as local property rights market failures remain ubiquitous, however, this paper makes the case that the global spatial allocation of water-intensive production plays a key role in mediating their welfare effects.

Fact 4: Agricultural policy plays a critical role in water use

An expansive literature in agricultural economics seeks to estimate the supply elasticity of crop production with respect to output prices, and more recent evidence establishes a direct link between agricultural policy and water consumption. For instance, on supply elasticities, Roberts and Schlenker (2013) estimate the impact of global prices on global production for four staple crops - rice, maize, soybeans, and wheat - and a number of other studies produce similar estimates for a single country (e.g. Hendricks, Smith and Sumner, 2014; Scott, 2014). The estimates showing that output prices meaningfully affect production imply an important role for subsidies, taxes, and tariffs that function by influencing prices. Some studies have gone further to demonstrate the role of policy directly, including papers such as Hendricks, Smith and Villoria (2018) that measure the impact of the Nominal Rate of Assistance (NRA) index described in Section 2 of this paper.

Estimates suggesting that agricultural policy plays an important role in driving crop production imply that policy also meaningfully affects water use, and previous work by Carleton (2021) quantifies this link directly. That paper measures the global impact of NRA on TWS using panel regressions that control for regional time fixed effects and a vector of time-varying climate controls. The analysis finds large effects of agricultural subsidies on water storage. On land with cropped area, each 10 percentage point increase in net agricultural subsidy changes the annual trend in total water storage by about 45 m^3/ha, approximately equal to moving from the median global parcel of arable land to the 25th or 75th percentile. A 100 percentage point movement in the NRA, which is well within the range of global policy variation in the cross-section and even over time within some individual countries, would take a location from the bottom 1% to the top 1% of global arable regions in TWS.

The effects of NRA on TWS in Carleton (2021) are larger for subsidies to more water-intensive crops and in regions with greater suitability for the subsidized crops, and are not detectable at all in placebo checks on regions without cropped area. Overall, the magnitude of the estimates suggests that even modest changes in agricultural policy can play a critical role in shaping the dynamic evolution of water availability over time. In contrast to the literature review described above in Fact 3 that suggests that local water market policies have had limited global reach, agricultural market interventions are ubiquitous around the world. The analysis in Carleton (2021) suggests that such policies can be a powerful tool to optimize global water consumption in agriculture. However, the degree to which existing policy, and the current distribution of agricultural activity, resembles the optimal allocation of global water resources remains undetermined.

Fact 5: Farmers largely plant water-intensive crops in water-abundant regions, but also in regions with rapidly declining water resources

If global water use were optimized in the production of tradable goods, we should expect water-intensive production to disproportionately occur in water-abundant locations. This prediction follows from standard trade theory, such as the Heckscher-Ohlin model, in which relative factor abundance drives comparative advantage. In our dynamic setting, places with more water resources can also sustain more water-intensive crops in the long-run without drawing down their stocks of water. Thus, while water is only one of many inputs that govern crop suitability (e.g. soil quality, temperature, precipitation) and we should not expect water availability to explain all the variation in global land use, an efficient allocation of global agriculture is likely to be broadly marked by water-intensive crops grown in water-abundant locations.

We evaluate this hypothesis by using the spatial datasets described in Section 2 to examine the global relationship between land use and water availability. We start by dividing the world's arable land into deciles of each of the three main measures of water resources: rainfall, groundwater table depth, and changes in total water storage (ΔTWS). For each water variable, we then calculate the proportion of each crop's global acreage that exists in each of the 10 deciles of water availability. Figure 2 shows the resulting patterns for a selection of major crops and water-intensive specialty crops.

The results in Figure 2 suggest that the correlation between planted acreage and water availability varies substantially across different water-intensive crops. For instance, rice, the most water-intensive of the major crops, is overwhelmingly concentrated in the world's most water-abundant regions. Over 93% of the world's rice acreage exists in locations with above median rainfall, and about 84% is in locations with above median groundwater tables (there is only a modest correlation between rainfall-abundant and groundwater-abundant regions, so these measures are largely independent). Even just the top decile of regions in both the rainfall and groundwater measures each account for a highly disproportionate share of global rice acreage around 20%. Overall, planting rice is ubiquitous in the world's wettest regions, but exceedingly rare in the world's driest regions.

Conversely, almonds, among the most water-intensive specialty crops, are primarily grown in relatively dry places. Over 85% of the world's almond acreage exists in locations below the global median of groundwater table depth, with well over half just in the bottom three deciles. Similarly, about 92% of almond-producing acreage is in locations below the median of global rainfall. Over 14% of the world's almonds grow just in California in regions below the world's 30th percentile of rainfall and the world's 15th percentile of groundwater tables. Given this discrepancy in patterns across individual crops, drawing broad conclusions about the spatial allocation of water-intensive agricultural activity requires a more comprehensive measure than analyzing crops one by one in isolation.

To produce a more systematic summary of the global relationship between crop choice and water resources, we construct a simple index of the water-intensity of land use across all crops. To do so, we start with the scientific estimates from Mekonnen and Hoekstra (2011) of the average global water-intensity of each crop, denominated in units m^3 of water used per acre planted for each crop. Using these estimates in conjunction with the Monfreda, Ramankutty and Foley (2008) data on acreage allocated to each crop, we calculate the acre-weighted average water-intensity of land use in m^3 per hectare for each parcel of global arable land. This calculation covers 115 crops that account for over 99.99% of total global planted acreage. Note that we assume that pasture land entails no human water use in this exercise, such that the measure of water-intensive land use also accounts for the extensive margin of whether to plant crops at all.

Figure 3 shows the average water-intensity of land use by decile of water availability for each of the three main variables. The results suggest that water-intensive agricultural production strongly clusters in those regions with high water tables and ample rainfall. The average water-intensity of land use increases nearly monotonically across deciles of both groundwater access and long-run average rainfall. On average, arable regions in the highest quintile of global groundwater tables have production that is more than twice as water-intensive as that of the lowest quintile. Even more strikingly, the rainiest quintile of places in the world use nearly five times more water per acre of arable land than the least rainy quintile.

This pattern of more water-intensive land use in wetter regions is driven both by allocating more arable land to crops, and by choosing more water-intensive crops conditional on planting. When restricting the analysis to cropped area, the top quintile of groundwater and rainfall regions choose crops that use about 20% and 200% more water per acre than the bottom quintile regions, respectively. Thus, these cross-sectional measures of land use and water availability suggest a strong role for resource abundance driving comparative advantage and patterns of specialization.

The correlation between water-intensive land use and the dynamic measure of water availability - trends in total water storage - shows a very different pattern in Figure 3. In contrast to the static measures of groundwater and rainfall in which water-intensive production is highly concentrated in water-abundant regions, Figure 3 shows that the decile of regions in the world losing water most rapidly in the period from 2003-2016 have the highest water-intensity of production. While the average acre of arable land in the world uses about 2500 m^3/ha of water, the average acre in the bottom decile of TWS uses over 3100 m^3/ha. Thus, while the first two panels of Figure 3 show that water-intensive production is concentrated in regions that are currently water-abundant, the third panel suggests that some locations planting water-intensive crops could be at risk of future depletion if current trends continue.

Examining Figure 3 further, it is worth also noting that the second through tenth global deciles of TWS show a modest positive correlation between trends in water availability and water-intensive land use, suggesting that the concurrence of water-intensive production and rapid water loss is restricted to a small number of regions. In these regions that comprise the bottom decile of TWS, the ratio of cropped acreage to arable land is about 21% higher than the world average (46% as compared to 38%), and the average water intensity of chosen crops conditional on planting is also about 25% higher than the world average. Rice, the most water-intensive staple crop, accounts for about 18% of cropped area and 27% of agricultural water consumption in these regions, compared with 16% and 45% for the world overall. Nearly 20% of the arable land area in the bottom decile of TWS is located in China, with about 10% each in India, Argentina, and Brazil.

Restricting attention to the 1% of global arable land losing water most rapidly reveals an even sharper pattern of water-intensive production. The average water intensity of land use in the bottom 1% is nearly 3700 m^3/ha, more than 10% higher than the bottom decile as a whole and almost 60% higher than the world average of 2500 m^3/ha. The proportion of cropped acreage is 58% higher in the bottom 1% than the global average, and rice accounts for about 20% of cropped area. About 75% of these extremely rapidly depleting regions are in three countries - India, Argentina, and Chile. To take a particularly extreme example, the Indian regions in the bottom 1% of TWS, which are primarily located near the capital of New Delhi and in the northeastern state of Uttarakhand, have crops planted on over 80% of their arable land with an average water-intensity of land use of over 8300 m^3/ha, well over three times the global average. These India regions are losing water at a rate of over 300 m^3/ha per year, a pace that will reduce the stock of water by the equivalent of annual average rainfall approximately every two decades.

While some of the regions losing water rapidly and planting water-intensive crops might not face a risk of serious water depletion in the near future, we also find some evidence of water-intensive land use in regions with declining resources and low existing stocks of groundwater in the cross-sectional data. Figure 4 shows the water intensity of land use for regions that fall in the bottom quartile of two or more of the three measures of water availability. Encouragingly, the figure shows that the most water-intensive land use occurs in regions that do not fall in the bottom quartile of any of rainfall, groundwater, or TWS, and that regions with low rainfall and low or declining stocks have very low water-intensity, consistent with the overall results in Figure 3.

Less favorably, the results in Figure 4 also show that the water-intensity of land use is substantial in those regions that fall in the bottom quartile of both the stock of groundwater and the trend in total water storage, which comprise about 7.1% of the world's arable land. Conditional on planting crops, the water-intensity of cropped area in those regions is about 20% higher than the global average, approximately equal to those regions that fall in the bottom quartile of none of the three variables. While the overall water-intensity is lower than the global average because only 26% of arable land is cropped, compared with 40% for the world as a whole, the data is consistent with a nontrivial role for water-intensive agricultural production in those regions with low and declining stocks of water that might be at the greatest risk of incurring serious costs from near term depletion.

Overall, the static measures of water availability in the data demonstrate a dominant pattern of water-intensive global agricultural production concentrating in presently water-abundant regions, but the dynamic measure of water trends reveals a small share of regions that have very water-intensive land use and rapidly declining water resources. The findings that the spatial pattern of global agricultural activity reflects an important role for water as a source of comparative advantage are consistent with those of Debaere (2014), which analyzed country-level patterns of exported products and water abundance. This paper is the first to collect the comprehensive global spatial data necessary to examine the relationship between agricultural activity and several measures of water availability across each granular parcel of arable land on Earth. In particular, the novel scientific datasets on groundwater table levels and trends in total water storage convey information about both the static and dynamic relationship between water and agriculture that have not been previously documented in the spatial environmental economics literature. The following sections provide further investigation of the welfare and policy implications of these facts.

4 Model

4.1 Basic environment

Time is discrete and indexed by $t \in \mathbb{N}$. The world economy consists of multiple countries, indexed by $i \in \mathcal{I} \equiv \{1, \ldots, I\}$, where consumption and production take place, and multiple aquifers, indexed by $q \in \mathcal{Q} \equiv \{1, \ldots, Q\}$, from which groundwater is extracted to be used in production.[16] Motivated by Fact 3, we make the simplifying assumption that the groundwater in each aquifer is a common-pool resource, available at the cost of extraction to any that cultivate the land above that aquifer.

Land is divided into heterogeneous fields, indexed by $f \in \mathcal{F} \equiv \{1, \ldots, F\}$, each of which is within some country i and above some aquifer q.[17] Fields comprise a continuum of heterogeneous parcels, indexed by ℓ. All fields correspond to 5-arc-minute grid cells. Because the surface of the earth is curved, grid cells at different latitudes cover different areas (with larger grid cells closer to the equator). We let h^f denote the area in hectares of field f.

Atomistic laborers in each country can choose to either farm their assigned parcel, earning the revenue from their harvest, or work for a wage w_i producing an outside good, which we think of as a composite of manufactured goods.[18] A farmer

---

[16] Aquifers occur naturally and therefore need not be circumscribed by country borders.

[17] Accordingly, objects that vary at the field level will not have country- or aquifer-specific subscripts. When we need to refer to the country in which field f is located, we'll write $i(f)$. Likewise, we'll write $q(f)$ for the aquifer below field f.

[18] Because laborers are assigned to parcels one-to-one, we sometimes use ℓ as an index of laborers.

uses his own labor to extract groundwater, which he combines with his land and remaining time to produce one of multiple crops, indexed by $k \in K \ \{1; \dots; K\}$.

4.1.1 Preferences

In each country i there is a representative household who lives hand-to-mouth and derives utility in each period from consuming the outside good, C^{o}_{it}, and an agricultural composite, C_{it}:

$$U_{it} = C^{o}_{it} + {}_{i} \ln C_{it} : \tag{1}$$

Since the upper-level utility function in equation (1) is quasilinear, there are no income effects. The total demand for crops depends only on a country-specific demand shifter, ${}_{i} \ \ 0$.

The agricultural composite, C_{it}, depends on the consumption of each crop, C^{k}_{it}, which itself depends on the consumption of varieties from different origins, C^{k}_{jit}:

$$C_{it} = \left[\sum_{k \in K} {}^{k}_{i}{}^{1/} \left(C^{k}_{it} \right)^{\frac{-1}{}} \right]^{\frac{}{-1}} \tag{2}$$

and

$$C^{k}_{it} = \left[\sum_{j \in I} {}^{k}_{ji}{}^{1/} \left(C^{k}_{jit} \right)^{\frac{-1}{}} \right]^{\frac{}{-1}} ; \tag{3}$$

where $ > 0$ denotes the elasticity of substitution between different crops (e.g., corn vs. soybeans) and $ > 0$ denotes the elasticity of substitution between different varieties of a given crop (e.g., Chinese vs. American soybeans). The last preference parameters, ${}^{k}_{i} \ \ 0$ and ${}^{k}_{ji} \ \ 0$, are crop- and crop-origin-specific demand shifters for country i.

4.1.2 Technology

In the agricultural sector, we assume that water and parcels of land are used in fixed proportions in the production of each crop. By combining land $L^{fk}_{t}(\omega)$ with groundwater $G^{fk}_{t}(\omega)$ and a fraction $H^{fk}_{t}(\omega)$ of his labor endowment, the farmer of parcel ω can produce

$$Q^{fk}_{t}(\omega) = A^{fk}(\omega) \left[H^{fk}_{t}(\omega) \right]^{} \left[\min \left\{ L^{fk}_{t}(\omega); \frac{G^{fk}_{t}(\omega)}{{}^{k}} \right\} \right]^{1-} ; \tag{4}$$

where $A^{fk}(\omega) \ \ 0$ denotes the total factor productivity of parcel ω in field f if allocated to crop k, and ^{k} measures the water intensity required to grow crop k anywhere.

The technology in (4) exhibits constant returns to scale in the three inputs, so the restriction to parcel-sized farms is without loss. Unlike the familiar Cobb-Douglas technology, which would impose a unit elasticity of substitution between each input, the technology in (4) imposes that land and groundwater be strict complements. This complementarity is intuitive: too little water and the crops will dry out and wilt; too much, and they will drown.[19]

Analogous to the approach in Eaton and Kortum (2002), we assume that TFP is independently drawn for each parcel ω from a field-specific Fréchet distribution:

$$P\left\{A^{f1}(\omega)\quad a^1;\dots;A^{fK}(\omega)\quad a^K\right\} = \exp\left\{\quad\sum_{k\in K}\left(\frac{a^k}{A^{fk}}\right)\right\} \tag{5}$$

where $\ > 1$ measures the extent of technological heterogeneity within each field and the constant $\ $ is set such that $A^{fk} = E[A^{fk}(\omega)]$.[20] The term $A^{fk}\quad 0$ measures the comparative and absolute advantage of a field in producing particular crops.

To extract the groundwater needed to irrigate his parcel ω, a farmer must allocate the remaining fraction $1\quad H_t^{fk}(\omega)$ of his labor endowment. Groundwater extraction is under constant returns to scale in the farmer's labor only. His productivity in extraction, however, is assumed to vary with the current depth of the water table below his parcel. Intuitively, it requires more labor to draw up one cubic meter of water from an aquifer with a low water table than from one whose table is near the surface. In particular, let D_{qt} denote the depth of the water table in aquifer q in period t. Then the corresponding labor productivity of groundwater extraction is given by

$$A_q^w(D_{qt}) = \ _q D_{qt}\ : \tag{6}$$

We recognize that groundwater is not the only source of water available to support crop production, which can also benefit from rainfall and surface water irrigation, and that the costs of groundwater extraction can vary across locations for a variety of reasons, such as technology or energy costs. Thus, we allow the parameter $_q$ to vary by aquifer, such that extraction productivity conditional on groundwater depth can differ across locations with differential access to rainfall, surface water, and technology, among other things.

The outside good is produced under constant returns to scale using labor only. The productivity of each worker in the outside sector, $A_i^o(\omega)$, is also drawn independently from a Fréchet distribution with the same shape parameter $\ $:

$$P\{A_i^o(\omega)\quad a^o\} = \exp\left\{\quad\left(\frac{a^o}{A_i^o}\right)\right\} \tag{7}$$

---

[19] Making the complementarity strict is for convenience. The main limitation of this specification is that it precludes farmers from adjusting how much water they use per unit of land. Appendix A.1.5 discusses a nested CES technology that generalizes that in (4).

[20] Formally, we set $\quad[(\quad 1)/\ ]\quad$, where $(\ )$ denotes the gamma function.

where $A^o_i = E[A^o_i(!\ )]$ is the average labor productivity in country i's outside sector.[21] Importantly, draws from this distribution are independent of the crop-specific productivity draws for the parcel the worker would otherwise till.

4.1.3 Market structure and trade costs

All markets are perfectly competitive. The outside good is freely traded and is used as the numeraire. International trade in crops, on the other hand, is subject to iceberg trade costs: In order to sell one unit of crop k to the representative consumer in country j, farmers in country i must ship ${}^k_{ij}$ units. The usual no-arbitrage condition then requires that the undistorted price of crop k produced in country i and sold in country j be equal to

$$p^k_{ijt} = {}^k_{ij}\, p^k_{it}; \tag{8}$$

where p^k_{it} denotes the local price of the domestic variety of crop k in country i before any taxes or subsidies are levied.

In addition to trade costs, crops are subject to policy distortions. Each national government sets a proportional tax ${}^k_{it}$ for each crop k at each date t.[22] Accordingly, the distorted farm gate price that farmers receive is ${}^k_{it}p^k_{it}$, and the distorted price paid by the consumer in j is ${}^k_{it}p^k_{ijt} = {}^k_{ij}\ {}^k_{it}p^k_{it}$. Agricultural policies are funded by lump-sum taxes on the domestic consumer such that the government's budget is always balanced.

The depth of the water table in aquifer q follows the law of motion

$$D_{qt+1} = D_{qt} + {}_q[(1 \qquad)X_{qt} \quad R_q] \tag{9}$$

where X_{qt} is the total amount of groundwater extracted from aquifer q in period t, R_q is the natural recharge of the aquifer (from rainfall, among other hydrological mechanisms), is the constant rate of return flow,[23] and ${}_q$ is an aquifer-specific conversion factor between volume and depth that depends on the local soil type.

4.2 Competitive equilibrium

In a competitive equilibrium, all consumers maximize their utility, all laborers maximize their returns, either by cultivating the revenue-maximizing crop on their parcel or by working in the outside sector, and all markets clear in each period.

---

[21] Some readers may find it odd that workers have heterogeneous productivities in the outside sector but not in the agricultural sector. Notice, however, that this model is observationally equivalent to one in which it is the workers *themselves* that vary in farming productivity across parcels within a field. The key assumption is then that the labor productivity of a worker is uncorrelated across tasks.

[22] If ${}^k_{it} < 1$, the policy is a net tax on that commodity; if ${}^k_{it} > 1$, it is a net subsidy.

[23] When water is poured onto a crop, only a fraction of that water is actually absorbed by the plant. The rest soaks back into the ground and, ultimately, back into the aquifer. The parameter accounts for the latter fraction.

4.2.1 Utility maximization

Given equations (1), (2), (3), and (8), utility maximization by the representative household in each country requires that[24]

$$C^k_{jit} = \frac{{}_i\,{}^k_i\,P^{k\;\;1}_{it}}{\sum_{\ell\in K}{}^\ell_i\,P^{\ell\;\;1}_{it}}\;\frac{{}^k_{ji}\;{}^k_{ji}\;{}^k_{jt}\,p^k_{jt}}{\sum_{n\in I}{}^k_{ni}\;{}^k_{ni}\;{}^k_{nt}\,p^{k\;\;1}_{nt}} \quad \text{for all } i; j \in I\,;\; k \in K\,; \quad (10)$$

where

$$P^k_{it} = \left[\sum_{n\in I}{}^k_{ni}\;{}^k_{ni}\;{}^k_{nt}\,p^{k\;\;1}_{nt}\right]^{\frac{1}{1}}$$

denotes the CES price index associated with crop k in country i at time t.

4.2.2 Revenue maximization and labor choice

Each laborer chooses to either cultivate his parcel or work in the outside sector. In the outside sector, profit maximization requires that workers are paid their marginal products whenever the outside good is produced. Throughout this paper, we assume that labor endowments are large enough and expenditure shares on agricultural goods are low enough for the outside good to be produced in all countries.

Should he choose to cultivate his parcel, the farmer must plant the crop that maximizes his revenue. Thus, given equation (4), land allocation can be solved as a simple discrete choice problem. To see how, first note that we choose land units such that $L^{fk}_t(\omega) = 1$ for each ω. The farmer allocates his labor optimally between drawing up water and tending the crop. One can show that his optimal output can be written as $A^{fk}(\omega)M({}^k; D_{q(f)t})$, where M is decreasing in the water intensity ^{k} and the current water table depth $D_{q(f)t}$.[25] Thus, the farmer's revenue from growing crop ℓ is

$$r^{f\ell}_t(\omega) = {}^{\ell}_{i(f)t}\,p^{\ell}_{i(f)t}\,A^{f\ell}(\omega)M({}^k; D_{q(f)t}):$$

We assume that laborers know their outside labor productivity and their vectors of crop-specific productivities before they select in which sector to work in each period. It follows that a laborer assigned to parcel ω in field f cultivates crop k on his parcel in period t with probability

$${}^{fk}_t \quad P\left\{r^{fk}_t(\omega) = \max\{A^o_{i(f)}(\omega); r^{f1}_t(\omega); :::; r^{fK}_t(\omega)\}\right\}:$$

---

[24] See A.1.1 for derivation.

[25] See the definition and derivation of M in A.1.2. An important point to note is that $\lim_{D\to\infty} M({}^k; D) = 0$ for any crop, so a depleted water supply implies an inability to produce crops.

Since there is a continuum of parcels within each field, π_t^{fk} also corresponds to the share of parcels on which crop k is cultivated in field f in period t. Because TFP and labor productivity in the outside sector are both independently distributed Fréchet with a common shape parameter according to equations (5) and (7), standard algebra implies that for all $f \in \mathcal{F}$ and $t \in \mathbb{N}$,[26]

$$\pi_t^{fk} = \frac{\left(\tau_{i(f)t}^k p_{i(f)t}^k A^{fk} M(\alpha^k; D_{q(f)t})\right)^\theta}{\left(A_{i(f)}^o\right)^\theta + \Phi_t^f}; \tag{11}$$

where

$$\Phi_t^f = \sum_{\ell \in \mathcal{K}} \left(\tau_{i(f)t}^\ell p_{i(f)t}^\ell A^{f\ell} M(\alpha^\ell; D_{q(f)t})\right)^\theta$$

summarizes the profitability of cultivating field f at time t. Looking at the expression in (11), one sees that the higher a crop's farm-gate price, $\tau_{i(f)t}^k p_{i(f)t}^k$, or mean productivity, A^{fk}, the higher the share of a given field allocated to that crop. The higher the crop-specific water requirements, α^k, however, the lower the share of that field allocated to crop k. The lower the labor productivity in water extraction due to a low water table, $D_{q(f)t}$, or the higher the mean labor productivity in the outside sector, $A_{i(f)}^o$, the lower the share of a given field allocated to any crops, as laborers are more likely to leave their parcels fallow to work in the outside sector.[27] Finally, the bigger the shape parameter θ, the less heterogeneity there is across parcels within a field, so the more sensitive farmers are to cross-crop differences in prices or average productivity.

Let $\mathcal{F}_i = \{f : i(f) = i\}$ denote the set of all fields in country i so that

$$Q_{it}^k = \sum_{f \in \mathcal{F}_i} \int_0^{h^f} Q_t^{fk}(\omega)\, d\omega$$

denotes the total output of crop k in country i. By equation (4) and the law of iterated expectations, it must be that

$$Q_{it}^k = \sum_{f \in \mathcal{F}_i} h^f \pi_t^{fk} M(\alpha^k; D_{q(f)t}) E\left[A^{fk}(\omega) \mid r_t^{fk}(\omega) = \max\{A_{i(f)}^o(\omega); r_t^{f1}(\omega); \ldots; r_t^{fK}(\omega)\}\right]:$$

---

[26] See A.1.3 for derivation.

[27] The share of laborers assigned to parcels in field f that choose to work in the outside sector at date t is given by

$$\pi_t^{fo} = \frac{\left(A_{i(f)}^o\right)^\theta}{\left(A_{i(f)}^o\right)^\theta + \Phi_t^f}:$$

Note, too, that so long as any crops are grown on a given field (which will happen so long as $D_q < \infty$), all crops with positive productivity, $A^{fk} > 0$, will be grown on that field.

Given our distributional assumptions, one can also check that[28]

$$E\left[A^{fk}(\omega)\, r_t^{fk}(\omega) = \max\{A^{o}_{i(f)}(\omega), r_t^{f1}(\omega), \ldots, r_t^{fK}(\omega)\}\right] = A^{fk} \quad {}_t^{fk} \quad {}^{1/} \;.$$

Note that because of the endogenous selection of fields into crops, the average productivity conditional on a crop being produced is strictly greater than the unconditional average. Combining the two previous expressions with equation (11), we obtain the following expression for the supply of crop k in country i:

$$Q_{it}^{k} = \sum_{f \in F_i} h^f A^{fk} M(\ {}^{k}, D_{q(f)t})\ \ {}_t^{fk} \ \ {}^{\frac{1}{\ }} \tag{12}$$

for all $i \in I$ and $k \in K$.

4.2.3 Market clearing and feasibility

Since trade in crops is subject to iceberg trade costs, market clearing for all varieties of all crops requires

$$Q_{it}^{k} = \sum_{j \in I} {}_{ij}^{k} C_{ijt}^{k} \quad \text{for all } i \in I \text{ and } k \in K. \tag{13}$$

Let $F_q = \{f : q(f) = q\}$ denote the set of all fields above aquifer q. Then total groundwater extracted from aquifer q in period t is

$$X_{qt} = \sum_{f \in F_q} \sum_{k \in K} h^f x_t^{fk} \ \ {}_t^{fk}. \tag{14}$$

where x_t^{fk} is optimal water extraction to grow crop k on field f.[29]

Finally, under the assumption that the outside good is produced in all countries, the amount of labor demanded by the outside sector adjusts to guarantee labor market clearing at a unit wage per efficiency unit of labor.

4.2.4 Definition and well-posedness of the competitive equilibrium

A competitive equilibrium in this environment is a feasible path, starting from an initial vector of groundwater depths, along which consumers maximize their utility, laborers maximize their returns, and all markets clear.

Definition 1. Given a set of agricultural policies, $\{\ {}_{it}^{k}\}$, and an initial vector of groundwater depths, $\{D_{q0}\}$, a competitive equilibrium is a path of consumption, $\{C_{jit}^{k}\}$, output, $\{Q_{it}^{k}\}$, undistorted prices, $\{p_{it}^{k}\}$, shares, $\{\ {}_t^{fk}\}$, groundwater depths,

---

[28] See A.1.4 for derivation.

[29] See the definition and derivation of x_t^{fk} in A.1.2.

$\{D_{qt}\}$, and groundwater extractions, $\{X_{qt}\}$, such that equations (9), (10), (11), (12), (13), and (14) hold.[30]

Conditions for the existence and uniqueness of an equilibrium are established in Appendix A.2. A key feature of the equilibrium is that it can be decomposed into a sequence of static sub-equilibria connected only through the law of motion (9). This is not because laborers are assumed to be myopic ad hoc; instead, it follows naturally from the fact that the groundwater in each aquifer is treated as a common-pool resource, the stock of which is always large relative to the farms that draw from it.[31] Accordingly, laborers do not consider how their choice of activity—and the water that must be extracted to do said activity—will affect the water table depth below their parcel in the future. Instead of solving a dynamic program, then, each laborer just solves a sequence of static problems. This is what makes estimation feasible at the fine spatial scales we have in our data, which we turn to next.[32]

5 Estimation

To simulate the model described in Section 4, we require estimates of a host of parameters governing the demand, supply, and hydrology blocks of the model. We cover each in turn. But, before doing so, we briefly revisit our data and describe the sample restrictions that we impose on the inputs to our estimation procedure. We conclude by checking the model's fit of targeted and untargeted moments.

5.1 Data and sample selection

We choose 2009 as the base year for our analysis because it is the midpoint of the period covered by the GRACE satellite data and the earliest year that postdates all groundwater table depths reported in Fan et al. (2013). It is also the most recent year for which data on agricultural output, land use, prices, and trade flows were all available from the FAOSTAT program at the FAO.

To construct our sample, we select 52 countries $(i \in I)$ that account for 94% of total GDP, 97% of total agricultural production value, and 99% of total agricultural

---

[30] There is nothing about the outside sector in this definition because that sector acts like a residual claimant on the resources of the economy once agricultural markets clear. See Appendix A.2.4 for discussion.

[31] That the representative consumers are assumed to be myopic is also crucial.

[32] Dynamic spatial models are notoriously hard to solve (Rossi-Hansberg, 2019). The framework developed in Desmet, Nagy and Rossi-Hansberg (2018) has recently been put to great use tackling environmental questions (see, e.g., Desmet, Kopp, Kulp, Nagy, Oppenheimer, Rossi-Hansberg and Strauss, 2021; Cruz and Rossi-Hansberg, 2021). Like ours, that framework relies on assumptions around agents' decision problems that ensure those problems are always static. In our case, the assumption is on the market structure (or lack thereof) for water; in their case, the assumption is on the returns to investment.

laborers (see Figure 5a).[33] Within these countries, we model land use at the resolution of gaez grid cells, restricting our attention to arable land, which leaves us with roughly 1.9 million fields. From the set of 38 crops for which gaez reports potential yields at the field level, we select 22 crops ($k \in K$) that account for 56% of total agricultural production and 59% of total water use. Our set of crops, shown in Figure 5c, includes major staples such as wheat, rice, maize, soybeans, and potatoes, a small number of water-intensive cash crops such as coffee and oil palm, and regional crops critical in many drier low-income regions, such as cassava, sorghum, millet, barley, and chickpeas. Importantly, the selected crops span a wide range of water intensities, from about 3,000 m^3/ha for yams to about 21,000 m^3/ha for bananas.

We construct "aquifers" ($q \in Q$) as clusters of contiguous grace pixels in the following way. First, following Richey et al. (2015), we overlay a map of the 37 largest global aquifer systems obtained from the Worldwide Hydrogeological Mapping and Assessment Program (WHYMAP) and group grace pixels into clusters that are at least partially contained within the borders of a single aquifer. Then, we overlay a map of 180 NASA-delineated water basins and group the remaining grace pixels into clusters according to the basin in which they lie. These basins are defined based on their hydrologic connectivity: within each basin, precipitation exits from the same location, so we take this as a useful measure of the geographic extent of the common pool resource externality. This procedure yields 278 clusters that partition global land area. Finally, we discard all grace pixels that lie fully outside the borders of our sample countries. The surviving 205 clusters map to what we call "aquifers" in the model.

5.2 Demand

To estimate the demand block, we closely follow the procedure described in Costinot, Donaldson and Smith (2016, §V.A). We proceed in three steps, moving outward from the innermost nest of the demand system in (10). First, we use data on prices, p_j^k, and bilateral trade flows by crop, E_{ji}^k, to estimate the elasticity of substitution between different varieties of a given crop and to invert a composite of the crop-origin-specific demand shifters, $^k_{ji}$, and trade costs, $^k_{ji}$.[34] Second, we use the previous estimates to construct price indices at the crop level, P_i^k, and combine them with data on crop expenditures, $E_i^k = \sum_{j \in I} E_{ji}^k$, to estimate the elasticity of substitution between crops and to invert the crop-specific demand shifters, ^k_i. Finally, we use

---

[33] Coverage shares are reported relative to totals for which data is available from faostat.

[34] Prices are measured in current US dollars per ton in 2009 from faostat. Where a US dollar price was not reported, we converted from local currency units using the exchange rates provided by the FAO. Where no price was reported, we followed Costinot, Donaldson and Smith (2016) in imputing the price from the fitted values of a regression of log prices on a country and crop fixed effect before imposing our sample restrictions. Trade flows are measured in current US dollars in 2009 from Comtrade via BACI. We compute autoconsumption, E_{ii}^k, as the difference between country i's total production and total exports of crop k.

data on total crop expenditures, $E_i = \sum_{k\in K} E_i^k$, to invert the uppermost demand shifters, α_i. Throughout, we will make repeated use of the identity

$$E_{ji}^k = \tau_{ji}^k p_j^k C_{ji}^k = \alpha_i \frac{\beta_i^k \left(P_i^k\right)^{1-\eta}}{\sum_{\ell\in K} \beta_i^\ell \left(P_i^\ell\right)^{1-\eta}} \frac{\beta_{ji}^k \left(\tau_{ji}^k p_j^k\right)^{1-\sigma}}{\sum_{n\in I} \beta_{ni}^k \left(\tau_{ni}^k p_n^k\right)^{1-\sigma}}; \tag{15}$$

which defines the value of exports of crop k from country j to country i.

Step 1: Estimating preferences across varieties of a given crop

With 52 countries and 22 crops, some crop-specific bilateral flows are zero. To rationalize these observations, whenever $E_{ji}^k = 0$ we set $\beta_{ji}^k(\tau_{ji}^k)^{1-\sigma} = 0$. Whenever $E_{ji}^k > 0$, on the other hand, we take logs and rearrange Equation (15) as

$$\ln\left(\frac{E_{ji}^k}{E_i^k}\right) = m_i^k + (1-\sigma)\ln p_j^k + \varepsilon_{ji}^k; \tag{16}$$

where the first term

$$m_i^k = -\ln\left[\sum_{n\in I} \beta_{ni}^k \left(\tau_{ni}^k p_n^k\right)^{1-\sigma}\right]$$

will be captured by an importer-crop fixed effect, and the last term $\varepsilon_{ji}^k \equiv \ln[\beta_{ji}^k(\tau_{ji}^k)^{1-\sigma}]$ is a structural error accounting for trade costs and unobserved variety-specific demand shifters. Without loss of generality, we normalize the demand shifters such that

$$\sum_{j\in I} \varepsilon_{ji}^k = 0. \tag{17}$$

By definition, the equilibrium prices are correlated with the structural errors. We instrument for prices using the log of the arithmetic average of the GAEZ potential yield of crop k across all fields in country j,

$$Z_j^k \equiv \ln\left(\frac{1}{|F_j|}\sum_{f\in F_j} A_j^{fk}\right).$$

The instrument Z_j^k should be correlated with undistorted crop prices, p_j^k, because higher productivity leads farmers in j to supply more of crop k. We assume that it is uncorrelated with the demand shifters and trade costs.[35]

From Equation (16), we estimate $\sigma = 5.32$ (with a standard error of 1.34 when clustered at the crop-importer and crop-exporter levels). The composite parameter $\beta_{ji}^k(\tau_{ji}^k)^{1-\sigma}$ is then backed out from the prediction error of the regression while imposing the normalization in Equation (17). This composite will be sufficient to construct equilibria—observed and counterfactual—in what follows.

---

[35] Formally, the exclusion restriction is $E[Z_j^k \varepsilon_{ji}^k] = 0$.

Step 2: Estimating preferences across crops

The second step will look much like the first, but instead of instrumenting for observed prices, we will need to instrument for a price index that we construct ourselves,

$$P_i^k = \left[\sum_{j \in I} \xi_{ji}^k (\tau_{ji}^k p_j^k)^{1-\sigma} \right]^{\frac{1}{1-\sigma}};$$

using the data and estimates from the previous step. With that index in hand, we rearrange Equation (15) as

$$\ln \frac{E_i^k}{E_i} = m_i + (1-\kappa) \ln P_i^k + \varepsilon_i^k \tag{18}$$

where the first term

$$m_i = -\ln \left[\sum_{\ell \in K} \beta_i^{\ell} (P_i^{\ell})^{1-\kappa} \right]$$

will be captured by an importer fixed effect, and the last term $\varepsilon_i^k \equiv \ln \beta_i^k$ is a structural error accounting for unobserved crop-specific demand shifters. Without loss of generality, we again normalize the demand shifters such that

$$\sum_{\ell \in K} \varepsilon_i^{\ell} = 0. \tag{19}$$

The same endogeneity concerns from the first step are present here, so we instrument for the price index, P_i^k, with the corresponding arithmetic average of potential yields for crop k, Z_i^k.

From Equation (18), we estimate $\kappa = 3.81$ (with a standard error of 0.29 when clustered at the importer level). As in the first step, the demand shifter β_i^k can then be inverted from the prediction error of the regression while imposing the normalization in Equation (19).

Step 3: Estimating preferences across sectors

The third step is the easiest: the utility function in (1) implies that $\beta_i = E_i$ for all $i \in I$. Across all three steps, our procedure for inverting the demand shifters has allowed us to match exactly the observed expenditures by each country on crops from each country in our sample in 2009.

5.3 Supply

To quantify the technologies in Equations (4)–(7), we first recall that the average potential yield of each field for each crop, A^{fk}, is directly observable in the gaez data.[36] Crucially, this is true for each field f regardless of whether field f is actually growing crop k. We calibrate the remaining parameters of the agricultural production function (4) to estimates from the literature. First, we set $1 \quad = 0.25$ to match the value-added share of land in agricultural production as computed in Boppart et al. (2019). Second, to quantify the crop-specific water intensities, ^{k}, we convert the estimates reported by Mekonnen and Hoekstra (2011) from cubic meters per ton to cubic meters per hectare using data on average yields (tons per hectare) from faostat.[37] Third, we set the elasticity of labor productivity with respect to depth in the pumping technology (6) to $\; = 1$ in order to accord with Burlig et al. (2021), who specify a cost function for groundwater extraction that is linear in the vertical distance over which the water must be lifted.

The remaining technological parameters that need to be estimated are the extent of within-field heterogeneity in potential yields, , each country's average labor productivity in the outside sector, $A^o \; \{A_i^o\}$, and each aquifer's scale of labor productivity in pumping groundwater, $\; \{\,_q\}$. Again inspired by Costinot, Donaldson and Smith (2016), our approach will be to choose values that allow us to best fit data on crop quantities and land use from faostat as well as water extraction implied by observed cropped area fractions from sage.

To that end, we define the predicted output of crop k in country i as a function of the parameters $(\,; A^o; \,)$ in a base year,

$$Q_i^k(\,; A^o; \,) = \sum_{f\in F_i} h^f A^{fk} M(\,^k; D_{q(f)}) \left(\frac{{}_i^k p_i^k A^{fk} M(\,^k; D_{q(f)})}{(A_i^o) \; + \sum_{\ell\in K} {}_i^\ell p_i^\ell A^{f\ell} M(\,^\ell; D_{q(f)})} \right)^{\frac{1}{}}$$

where the dependence on is through the optimal input bundle, M. Likewise, we define the predicted land allocated to agriculture in country i and the predicted water extracted from aquifer q as functions of the same parameters in a base year,

$$L_i(\,; A^o; \,) = \sum_{k\in K} \sum_{f\in F_i} h^f \frac{{}_i^k p_i^k A^{fk} M(\,^k; D_{q(f)})}{(A_i^o) \; + \sum_{\ell\in K} {}_i^\ell p_i^\ell A^{f\ell} M(\,^\ell; D_{q(f)})}$$

$$X_q(\,; A^o; \,) = \sum_{k\in K} \sum_{f\in F_q} h^f x^{fk} \frac{{}_i^k p_i^k A^{fk} M(\,^k; D_{q(f)})}{(A_i^o) \; + \sum_{\ell\in K} {}_{it}^\ell p_{it}^\ell A^{f\ell} M(\,^\ell; D_{q(f)t})} :$$

---

[36] We use the baseline, high-input gaez estimates, measured in dry-weight tons per hectare and converted to fresh-weight tons using the provided conversion table (in order to accord with faostat production data). gaez reports separate estimates for rainfed and irrigated fields, so we use observed shares of irrigated land from the Global Map of Irrigation Areas to compute a weighted average estimate. For rice, gaez reports two varieties (dryland rice and wetland rice), but faostat only reports the aggregate category, rice. In this case we use the maximum yield over the two varieties for each field.

[37] Note that we hold ^{k} fixed across crops throughout the world because we are not aware of any systematic estimates of how water-intensity for each crop differs across regions. Engineering estimates suggest that differences in water-intensity across irrigation methods within crop are small relative to differences between crops.

Note that all three functions encode an equilibrium for a given set of agricultural policies, ${}_i^k$. In the estimation, we set these to match the nominal rates of assistance (NRAs) reported for 2009 by the World Bank dai project.[38]

Conditional on , we will search for the vectors A^o and at which, simultaneously, (i) the total amount of land allocated to crops predicted by the model for each country i, $L_i(; A^o;)$, exactly matches the total amount of land allocated to crops per faostat , $\hat{L}_i$, expressed in hectares, and (ii) the total amount of water extraction predicted by the model for each aquifer q, $X_q(; A^o;)$, exactly matches the total amount of water extraction implied by observed cropped area fractions from sage, $\hat{}^{fk}$, and the calibrated water intensity of each crop,

$$\hat{X}_q = \sum_{k \in K} \sum_{f \in F_q} h^f \; {}^k \hat{}^{fk};$$

expressed in cubic meters. Given A^o and , we then search for the value of at which output predicted by the model, $Q_i^k(; A^o;)$, best matches observed output of crop k in country i per faostat , $\hat{Q}_i^k$, expressed in fresh-weight tons. Formally, we use nonlinear least squares (NLS) to estimate $(; A^o;)$ as the solution of

$$\min_{ ; A^o ; } \sum_{i \in I} \sum_{k \in K} \left[\ln Q_i^k(; A^o;) \quad \ln \hat{Q}_i^k \right]^2 \tag{20}$$

subject to

$$L_i(; A^o;) = \hat{L}_i \quad \forall i \in I; \tag{21}$$

$$X_q(; A^o;) = \hat{X}_q \quad \forall q \in Q. \tag{22}$$

In the current draft, we consider only a restricted version of the NLS procedure just described. For now we simply calibrate $= 2.46$ to match the analogous estimate from Costinot, Donaldson and Smith (2016).[39] We then search for the vectors A^o and that simultaneously minimize the mean squared deviations between the model predictions and the data for land use and water extraction.[40]

---

[38] Nominal rates of assistance are reported in percentage terms, so ${}_i^k = 1 +$ NRA${}_i^k$. NRAs are not reported for every country-crop pair in 2009. Where no NRA was reported, we set ${}_i^k$ to the "general" value that summarizes country i's assistance across agricultural products in 2009, if available. Where no general value was reported, we set ${}_i^k = 1$.

[39] Costinot, Donaldson and Smith (2016) use a bootstrap procedure, with replacement at the country level, with 400 replications to estimate the 95 percent confidence interval around their estimate, which they find to be $[2.28; 2.62]$. ?, using Peruvian data on crop output disaggregated at the district level, estimates $= 2.06$.

[40] Although we can, in principle, exactly satisfy the constraints (21)–(22), doing so requires adjusting the naive data targets $\hat{L}_i$ and $\hat{X}_q$ to something attainable by the model with the current menu of crops, as we describe more below.

5.4 Hydrology

We close the quantitative model by specifying the remaining hydrologic parameters that govern the law of motion (9). We assume a common return flow rate across crops and aquifers, $= 0.25$, which is the central value reported by Dewandel et al. (2008).[41] For each aquifer q, its initial depth, D_{q0}, is set to the median value across constituent grid cells as reported in Fan et al. (2013), expressed in centimeters from the surface. To convert from a change in volume in aquifer q (coming in part from extraction, X_{qt}) to a change in depth, $D_{qt+1} \; D_{qt}$, we calibrate the parameter ${}_q$, which accounts for two characteristics of the aquifer. First, and most obviously, its area, which can be computed easily from the underlying GRACE pixels and expressed in squared meters. Second, its specific yield, which is the volume of water that can be drained from porous media by gravity, relative to the total porous media volume. Specific yield depends on soil type, which itself depends on location and depth.[42] Accordingly, we use maps of soil type from Hengl et al. (2017) and estimates of specific yield by soil type from Loheide et al. (2005) to calibrate ${}_q$.

The final component is the natural recharge rate for each aquifer q, R_q, expressed in cubic meters. We calibrate recharge in order to match the average annual change in total water volume for each aquifer implied by GRACE trends from 2003 to 2016, given what the model implies about water extraction. This captures the dominant source of recharge—rainfall—but also accounts for runoff between aquifers or any other unobserved variation in the global hydrological system.

5.5 Goodness of fit

We compare the fit of the model to the data along the most important dimensions for the simulations that follow: cropped area, which we target at the country level, and water extraction, which we target at the aquifer level. Figures 6a and 6b show the percent difference between the model simulation and the data for cropped area and water extraction, respectively. For cropped area, the model reproduces the global patterns almost exactly for the great majority of the world, with a simulation error under 5% for 43 out of 52 countries that cover 71% of the total cropland. Similarly for water extraction, the simulation error is under 5% for 145 of the 205 aquifers that cover 67% of the arable land represented in the model. Recall that we are attempting to match all agricultural water use for 115 crops contained in the Monfreda, Ramankutty and Foley (2008) data on cropped area and the Mekonnen and Hoekstra (2011) estimates of water-intensity by crop, but there are only 22 crops in the model simulations, which are constrained by the availability of GAEZ potential yields data, so we do not expect the model to be able to fit perfectly everywhere.

---

[41] In future versions of the paper, we will account for heterogeneous return flow rates across crops.

[42] For example, clay is among the least porous types of soil; gravel, among the most porous.

The geographic pattern of the simulation errors shown in Figures 6a and 6b makes apparent that the model's misses are heavily concentrated in southeast Asia, where the simulations consistently underpredict both cropped area and agricultural water extraction. There are several unique features of this region, which includes Indonesia, Malaysia, the Philippines, Thailand, Vietnam, and Bangladesh, including extremely high rates of rainfall totaling several times the global average (Figure 1a) and the widespread availability of both surface water and easily accessible water tables close to the surface (Figure 1b).

Most notably, we emphasize that the aquifers in southeast Asian countries contain extreme outliers in the data on the water intensity of land use by farmers, as shown in Figure 6c. For most of the land in these regions, the water intensity in the data we target suggests that the average acre of arable land uses over 10,000 m^3=ha. For context, recall that rice uses 8,790 m^3=ha. Thus, matching the data would require the model to simulate arable land use in which farmers choose to crop 100% of acreage with a crop that uses more water than the most water-intensive staple crop, which is clearly infeasible in the model. At the time of this draft, our best understanding is that the extreme outliers in the land use targets for southeast Asia occur because many farms in these regions crop the same acreage multiple times in the same year. The data we use contains estimates of harvested acreage, but while the same acre can be harvested more than once in some locations in practice, it only exists once in each year in the model. So it is possible that the model's targets are actually impossible to meet in some locations within this region.

In future drafts of the paper, we plan to make adjustments for how the model represents potential land use in regions with multiple cropping cycles in each year by using global spatial datasets on crop calendars that identify these locations. For now, we emphasize the map in Figure 6d, which shows that the model reproduces the qualitative pattern of unusually extreme water-intensity of land use in these countries in southeast Asia even if it does not fully capture the magnitudes of their agricultural water consumption.

As an additional assessment of how well the model captures patterns of global water resources, we also examine the calibrated values of the extraction productivity parameters across aquifers. Recall from Equation (6) that the parameter governs the labor productivity of water extraction for farmers, conditional on aquifer depth. We allow this parameter to vary across aquifers because the costs of extraction, conditional on the water table level, can vary across regions for a number of reasons. These include the free availability of water through rainfall, the potential for surface water irrigation that could be less expensive, heterogeneous irrigation technologies available to farmers, and the local institutions that govern water access. While not all the determinants of the input costs of water are well measured in the data, we might expect the values to be correlated with plausibly relevant observable covariates.

Figure 7 shows raw correlations of with other aquifer characteristics. Encouragingly, the top-left panel and top right-panel show that, over most of the range of values represented in the data, both precipitation and the share of area equipped for irrigation are positively correlated with . A few regions with more extreme precipitation values have lower extraction productivities, consistent with the agronomic literature that suggests a U-shaped relationship between moisture and crop productivity. The bottom-left panel also shows that regions with greater nighttime lights, a proxy for local income, have higher extraction productivities, consistent with a role for technology. The bottom middle panel shows no relationship between and our measure of local surface water presence, and the top middle panel shows a U-shaped relationship between temperature and , which is again consistent with relevant estimates in agronomy and agricultural economics.

The covariates shown in Figure 7 are correlated with each other in various ways, so Table 1 also shows their partial correlations with the effective input costs of water. The first column shows a regression in which the dependent variable is the log of , which governs the extraction productivity for the farmer for a given level of water table depth, and the second column shows results for the log of the extraction productivity, $A^{w}_{q}(D_{qt})$, that also incorporates the actual water table depth in the initial period. The patterns in the regression are very similar to those of the raw correlations in Figure 7. The data clearly reveal familiar U-shaped patterns in precipitation and temperature that imply an optimal level of rainfall and temperature for water-intensive agriculture. The results also show a positive role for area equipped with irrigation and nighttime lights in promoting extraction productivity. The only puzzling result is the lack of relationship between extraction productivities and the Pekel et al. (2016) data on surface water presence.

Perhaps most notably, the final row of the table shows that extraction costs vary strongly with groundwater table depth. Recall from Section 5.4 that we infer to match the data on the water-intensity of land use. Recall also the patterns shown in Figure 3 and Fact 5 of Section 3 that show that the most water-intensive global land use is in regions with higher water tables and more rainfall. Thus, it is the revealed preference of the world's farmers deciding how to allocate their land from which we infer the extraction productivities. Overall, the parameter estimates suggest a critical role for the natural environment in governing the input costs of water, which are much higher in dry locations. These estimates are important for the counterfactuals that follow both because they help explain the mechanism by which scarcity can map into comparative advantage and specialization for unpriced water inputs, and because they quantitatively determine how changes in the water table over time map into productivity and welfare.

6 Counterfactuals

6.1 Autarky

We start by addressing the overarching question of the general impacts of global agricultural trade on water depletion and long-run welfare. To do so, we consider what global agriculture and water resources would look like in the complete absence of international trade by running a counterfactual scenario in which we set the iceberg trade costs, ${}^{k}_{ij}$, to infinity for all countries and crops. We run each counterfactual forward for 30 years to show the dynamic evolution of water availability and welfare with and without international trade.

Before examining the spatial distribution of agriculture and water resources in a hypothetical world without agricultural trade, it is worth highlighting some overall global patterns. Figure 8 compares how cropped area, water extraction, aquifer depth, and welfare evolve in the baseline and in autarky over a 30-year period starting in the initial year of 2003. The first key result shown in panel (a) is that global cropped area is nearly 90% higher in autarky than in the baseline. This reflects the static efficiency gains from trade. When global agriculture can be produced far from where it is consumed, the most efficient agricultural regions of the world can disproportionately be used to feed the global population. In autarky, however, each country must produce enough to meet domestic consumption, even if doing so requires using less productive land. Thus, global rice and wheat yields are 15% and 4% lower on impact, respectively, in autarky than in the baseline.

Panel (b) of Figure 8 shows that the additional cropland required to meet global demand in autarky also raises global water extraction by about 60% in the initial period. We interpret this number as the aggregate water savings that allowing trade creates by improving the spatial efficiency of production and reducing the land used for global agriculture. Notably, the autarky increase in global water extraction is nearly 30 percentage points smaller than the increase in global cropped area. This is because the model simulations predict that some of the adjustment in response to countries having to produce all their food domestically would come through the margin of substituting to less water-intensive crops. Conditional on cropping, the water intensity of land use is about 15% lower in autarky than in the baseline in the initial period, as lower water-intensity crops grow as a share of the global distribution of calories in the absence of the spatial efficiency gains from trade.

Panel (c) of Figure 8 shows the evolution of average global aquifer depth in the baseline and in autarky. As discussed in Section 3, global arable land in the data is evenly split between locations gaining and losing water. We match these existing trends in water availability exactly in the model calibration, so the blue line in Figure 8c shows that the baseline mean change in aquifer depth across global arable land is essentially zero over time. In contrast, global average aquifer depth falls by over 5 meters after 30 years of autarky, meaning water tables decline by about 27% of their initial average values that we take from the Fan, Li and Miguez-Macho

(2013) data. Without trade, much more land and much more water is required to meet global demand for food, which rapidly depletes the world's water resources.

Falling water tables in autarky have important implications for the evolution of global agriculture and for welfare. Panel (d) shows that global welfare is stable over time in the baseline counterfactual, as there is no average trend in water tables. In autarky, however, welfare starts falling immediately upon impact, and is 2.2% lower for the world after 30 years. The graph normalizes welfare across the counterfactuals in the initial period in order to net out the static gains from trade, so the 2.2% number measures only the dynamic welfare gains from trade that come through the channel of preserving the world's water resources over time. This is a large effect in the context of previous estimates of static welfare gains in the trade literature (Arkolakis, Costinot and Rodrguez-Clare, 2012), though we caution that the welfare estimates in this draft of the paper remain preliminary.

Declining water tables affect welfare through several channels in the model. When water becomes more difficult to access, agricultural productivity falls, food prices rise, and production in the outside sector declines as more land and labor are allocated to producing enough food to meet demand, as shown in Figure 8a. Global average yields of rice and wheat, which are already 15% and 4% lower in the first period of autarky due to the less efficient allocation of global land, fall by an additional 8% and 3% after 30 years of water depletion. In addition, the composition of consumption also shifts away from water-intensive crops. Panel (b) of Figure 8 shows that global water extraction falls substantially over time in autarky as water tables fall and farmers respond accordingly to the rising cost of obtaining water. The average water-intensity of cropped land starts 15% lower in autarky as farmers in dry regions have to grow more food for domestic consumption, and falls further to 30% lower after 30 years as farmers substitute even more away from water-intensive crops over time. Feeding the world's population in the absence of trade requires substituting somewhat away from water-intensive consumption at first, and even moreso over time as water tables decline.

The spatial distribution of changes in autarky shows that the rise in cropped area and water extraction has a long right tail that skews heavily towards dry regions of the world. While extraction rises by 60% upon impact for the world overall, Figure 9a shows that it increases by as much as 300-600% in a number of regions, including most of Africa and Australia. As the map implies, the regions in which cropped area and water extraction increase most tend to be drier than the global average. Figure 9b shows that the increase in water extraction is over 150% for regions in the bottom quartile of initial water table depth, and nearly 200% for regions in the bottom quartile of rainfall, compared to 60% for the world on average. In contrast, extraction increases by only about 20-25% in regions in the top quartile of initial depth or rainfall. Some extremely wet food-exporting regions in Southeast Asia with very water-intensive production in the baseline even see declines in their water consumption in autarky. The dominant pattern in the baseline equilibrium, as reflected in Fact 5 from Section 3, is that wetter regions in the world export water-intensive goods to drier regions, so the scenario where all food is produced domestically forces water consumption to rise most in the predominantly importing regions that can least support it.

Figure 10 shows changes in water table depth across the world in the baseline and in autarky. Recall from Section 5 that the change in depth is a function of not only water extraction and recharge (which is fixed in each aquifer across years and counterfactuals), but also of the specific yield, , in each aquifer, which maps changes in volume to changes in depth based on local soil characteristics. In the baseline, the arable regions of the world that are depleting most rapidly, such as northwestern India and California's Central Valley, project to have their water tables decline by about one to three meters over the 30 year horizon. In autarky, the average global decline is about five meters. For the driest regions at the 90th percentile of water table depth, the 30-year decline is an enormous 18 meters of depth. Overall, about 16% of the world's arable land sees their water tables fall by over 10 meters within thirty years in autarky, a rate of water depletion that far exceeds anything observed in the existing data.

Figure 11 shows the evolution of groundwater table depth and welfare over time in two selected example countries, Australia and India. Australia is a major food-importing region in the baseline, preserving its domestic water resources by largely relying on other countries for their water-intensive consumption. On average, its aquifers are not depleting in the baseline, as shown in the blue line of panel (a). In autarky, however, the red line shows that Australia's water tables decline substantially. Panel (b) shows that welfare falls correspondingly over time in autarky, whereas it holds steady in the baseline. The decline in welfare in Australia is smaller than in some other countries because the agricultural share of GDP is smaller, though note that the welfare calculations in the paper remain preliminary.

Panel (c) of Figure 11 shows that aquifer depth in the northwestern region of India is depleting substantially in the baseline simulation, but is stable in autarky. Patterns elsewhere in India are qualitatively similar, such that panel (d) shows that welfare declines notably in India in the baseline, and declines somewhat less over time in autarky. This is consistent with the evidence in Sekhri (2022) that agricultural trade exacerbates depletion in India - a major food exporter. Note as a caveat to this result that India is one of the small proportion of regions in the world for which we do not closely match baseline water consumption (see Figure 6). However, since we underpredict water consumption in India in the simulation relative to the data, we expect that improving the fit of the model would only strengthen this pattern of agricultural trade exacerbating, rather than preventing, depletion.

While it is difficult to see visually, the maps in Figures 9a and 10b also show that the Central Valley of California constitutes a sharp exception to the broad global pattern of trade preventing water depletion. Much like northern India, this region of California with export-oriented agriculture and rapid depletion in the baseline has lower water extraction and a stable trend in its water table in autarky. We take these counterexamples to the global patterns as an encouraging sign that the model's calibration and simulation does not force a result in which trade liberalization prevents water depletion. In fact, the results show that the two prominent examples in the literature of possible trade-induced local water depletion appear to be consistent with a dynamic general equilibrium analysis. However, the results also show that California and India stand in sharp contrast to the overall picture across the world, in which agricultural trade overwhelmingly serves to prevent rather than create problems associated with water depletion.

6.2 Uruguay Round Agricultural Trade Policy Reforms

We now turn from analyzing the general impacts of global trade on water resource depletion and welfare to considering specific realized and potential policy reforms. We start by considering the set of trade agreements arising from the Uruguay Round of World Trade Organization negotiations that concluded in 1994. These agreements were signed by 123 nations, and together marked the most substantial liberalization of global agricultural markets observed in history. Prior to the Uruguay Round, agriculture was largely omitted from the General Agreement on Tariffs and Trade (GATT) and national governments were heavily leveraging subsidies, import restrictions, and foreign exchange market manipulations, among other tools, to distort agricultural production and trade (Healy et al., 1998). The agreements these countries signed in 1994 induced a sweeping array of reforms that reduced global agricultural market distortions in most countries.

To investigate the implications of this observed market liberalization on water depletion and welfare, we run a counterfactual simulation in which we fix all country-by-crop Nominal Rates of Assistance (NRAs; $_i^k$ from Section 4) to equal their average values during the period from 1986-1994, the negotiating years of the Uruguay Round. Comparing this counterfactual to the baseline in which we use NRA values observed in 2009 allows us to contrast the long-run effects of the policy regimes that existed before and after the reforms were made. In addition, because many production and trade distortions remain in effect today, we also conduct a simulation in which we set all NRA values to zero. This allows us to evaluate the effects of removing all remaining policy distortions in global agricultural markets.

Figures 12a and 12b show that these two alternative scenarios lead to very distinct spatial distributions of reforms. Before the Uruguay Round, wealthy nations were heavily subsidizing agricultural production, while developing economies relied on import-substitution policies intended to promote industrialization, including agricultural export taxes and foreign exchange manipulations that suppressed agricultural production (Anderson et al., 2013). Therefore, the Uruguay Round reforms caused NRAs to decline in regions like the United States, Europe, and Australia, but increased NRAs in many lower and middle-income regions, including India, Brazil, and many parts of Africa (Figure 12a). In contrast, remaining distortions in agricultural policy consist almost entirely of net subsidies, such that removing these interventions would reduce production incentives nearly everywhere (Figure 12b).

The maps in Figure 12c and 12d show the simulated change in global water extraction by region resulting from the Uruguay Round reforms and the hypothetical elimination of remaining agricultural net subsidies, respectively. As expected from the pattern of agricultural policy changes, the Uruguay Round reforms reduced global water use in richer, largely wetter, regions such as western Europe and the U.S. that reduced their domestic supports to agriculture. In contrast, many developing countries increased extraction substantially as their reforms largely removed disincentives to agricultural production. On net, the simulations suggest that the Uruguay round of reforms actually increased global water extraction by about 5%. In largely water-scarce lower-income regions especially, the policy changes raised consumption relative to recharge, increasing the rate of depletion and causing declines in welfare over time. In contrast, Figure 12d shows that hypothetical further policy liberalizations that remove remaining agricultural distortions would reduce global extraction in most regions, by about 5% in total, helping preserve the world's water resources and mitigate any dynamic declines in welfare.

Together, these findings underscore the complex and nuanced relationship between agricultural trade, rates of water depletion, and welfare. While the autarky counterfactuals overwhelmingly point to substantial economic and environmental gains from trade along both static and dynamic resource conservation dimensions, we find that the most historic agricultural market liberalization to date appears to have exacerbated global water stress and reduced welfare over time. In contrast, further liberalizations that eliminate remaining distortions would broadly have the opposite effect. These results, taken together with the spatial patterns showing that autarky preserves water resources in a small number of rapidly depleting regions, such as California and northern India, make clear that trade in agricultural goods need not be globally beneficial for water or welfare. In a setting with ubiquitous input market failures, it is critical to assess the specific spatial arrangement of any alternative output policy regimes in order to determine their likely impacts on water resources and welfare.

7 Conclusion

This paper considers the impact of global agricultural trade policy on regional water scarcity and long-run welfare. In a setting in which input markets are notoriously distorted and largely do not exist, it is possible for trade to exacerbate the local common pool resource externality and cause dynamic losses from depletion over time if poor institutions act as a source of false comparative advantage (Chichilnisky, 1994).

However, we show that in the case of water, the relative physical scarcity of the resource across locations maps strongly into its effective input price, such that water-abundant regions do have strong comparative advantage in water-intensive production despite the widespread lack of functioning markets for the input. Thus, global agricultural trade allows specialization in water-intensive production to cluster in water-abundant regions, preventing severe global depletion of water resources in the vast majority of the world, and especially so in dry locations that presently rely heavily on food imports. The model allows us to introduce and quantify a previously unmeasured channel of the dynamic resource availability gains from trade in this context. Despite these broad beneficial effects of global trade for water resources, we also show that trade and agricultural policy liberalizations can have the harmful effect of exacerbating depletion in some locations and in the case of some observed and hypothetical policy reforms, suggesting that the relationship between trade policy and water resources is nuanced and case-specific.

A natural question is whether the findings about trade preventing depletion of water resources can be generalized to the cases of other open access natural resources. Given the mechanism for the primary results in this paper, we conjecture that a key factor determining the relationship between trade and resource depletion is whether there exists a physical mechanism for the resource's scarcity to map into its marginal cost of extraction. If so, it may be the case that nature can, at least partially, compensate for nonfunctional input markets in governing the allocation. If not, it may not be that trade is helpful at all - and indeed could be harmful - for preserving the long-run availability of the resource. Thus, we might expect to see very different results for the effects of global trade on the management of forests, for which the marginal cost of extraction is generally independent of the existing stock of forests, and the management of fisheries, for which the marginal cost of extraction depends strongly on the existing local stock of fish.

We conclude with a final suggestion for future research. One policy implication of the results in this paper is that international trade in agriculture creates considerable benefits through its allocation and preservation of water resources. Yet policymakers often oppose reforms that reduce barriers to importing food because of a belief that domestic production is necessary to enhance \food security." In this paper, we show an important dimension through which agricultural trade actually enhances this definition of food security. In the absence of imports from water-abundant regions, many countries in water-scarce locations would draw down their water resources substantially enough that domestic agricultural production would become much more difficult in the long-run, eventually increasing their reliance on external producers. In the context of these findings, and of the broader recent developments in the literature on trade, agriculture, and the environment, it may be worth considering more closely the relationship between trade policy and reliable long-run access to food.

References

Adamopoulos, Tasso and Diego Restuccia, \The size distribution of farms and international productivity differences," American Economic Review, 2014, 104 (6), 1667{97.

Allan, John A, \Virtual water: a strategic resource," Ground water, 1998, 36 (4), 545{547.

Allan, Tony, Virtual water: tackling the threat to our planet's most precious resource, Bloomsbury Publishing, 2011.

Allen, Treb, Costas Arkolakis, and Yuta Takahashi, \Universal gravity," Journal of Political Economy, 2020, 128 (2), 393{433.

Alvarez, Fernando and Robert E. Lucas Jr., \General equilibrium analysis of the Eaton{Kortum model of international trade," Journal of Monetary Economics, 2007, 54 (6), 1726{1768.

Anderson, Kym, Gordon Rausser, and Johan Swinnen, \Political economy of public policies: Insights from distortions to agricultural and food markets," Journal of Economic Literature, June 2013, 51 (2), 423{77.

—, Kym Anderson, Marianne Kurzweil, Will Martin, Damiano Sandri, and Ernesto Valenzuela, \Methodology for measuring distortions to agricultural incentives," Technical Report, World Bank 2008.

Arkolakis, Costas, Arnaud Costinot, and Andres Rodrguez-Clare, \New trade models, same old gains?," American Economic Review, 2012, 102 (1), 94{130.

Ayres, Andrew B., Kyle C. Meng, and Andrew J. Plantinga, \Do environmental markets improve on open access? Evidence from California groundwater rights," Journal of Political Economy, July 2021, 121 (10).

Boppart, Timo, Patrick Kiernan, Per Krusell, and Hannes Malmberg, \The macroeconomics of intensive agriculture," July 2019.

Bruno, Ellen M. and Katrina Jessoe, \Missing markets: Evidence on agricultural groundwater demand from volumetric pricing," Journal of Public Economics, April 2021, 196.

Burlig, Fiona, Louis Preonas, and Matt Woerman, \Energy, groundwater, and crop choice," Working Paper 28706 April 2021.

Carleton, Tamma, \The global water footprint of distortionary agricultural policy," July 2021. Presented at NBER SI 2021.

Carr, JA, David A Seekell, and P D'Odorico, “Inequality or injustice in water use for food?,” Environmental Research Letters, 2015, 10 (2), 024013.

Carr, Joel A, Paolo D'Odorico, Francesco Laio, and Luca Ridolfi, “Recent history and geography of virtual water trade,” PloS one, 2013, 8 (2), e55825.

Chichilnisky, Graciela, “North-south trade and the global environment,” American Economic Review, 1994, 84 (4), 851–874.

Copeland, Brian R., Joseph S. Shapiro, and M. Scott Taylor, “Globalization and the environment,” in Gita Gopinath, Elhanan Helpman, and Kenneth Rogoff, eds., Handbook of International Economics, Vol. 5, Elsevier, 2022, chapter 2, pp. 61–146.

Costinot, Arnaud, Dave Donaldson, and Cory Smith, “Evolving comparative advantage and the impact of climate change in agricultural markets: Evidence from 1.7 million fields around the world,” Journal of Political Economy, 2016, 124 (1), 205–248.

Cruz, José-Luis and Esteban Rossi-Hansberg, “The economic geography of global warming,” Working Paper 28466, National Bureau of Economic Research February 2021.

Dalin, Carole, Yoshihide Wada, Thomas Kastner, and Michael J Puma, “Groundwater depletion embedded in international food trade,” Nature, 2017, 543 (7647), 700–704.

Debaere, Peter, “The global economics of water: Is water a source of comparative advantage?,” American Economic Journal: Applied Economics, April 2014, 6 (2), 32–48.

——, Brian D Richter, Kyle Frankel Davis, Melissa S Duvall, Jessica Ann Gephart, Clark E O'Bannon, Carolyn Pelnik, Emily Maynard Powell, and Tyler William Smith, “Water markets as a response to scarcity,” Water Policy, 2014, 16 (4), 625–649.

Desmet, Klaus, Dávid Kristián Nagy, and Esteban Rossi-Hansberg, “The geography of development,” Journal of Political Economy, 2018, 126 (3), 903–983.

——, Robert E. Kopp, Scott A. Kulp, Dávid Krisztián Nagy, Michael Oppenheimer, Esteban Rossi-Hansberg, and Benjamin H. Strauss, “Evaluating the economic cost of coastal flooding,” American Economic Journal: Macroeconomics, April 2021, 13 (2), 444–486.

Dewandel, B., J.-M. Gandolfi, D. de Condappa, and S. Ahmed, “An efficient methodology for estimating irrigation return flow coefficients of irrigated crops at watershed and seasonal scale," *Hydrological Processes*, 2008, *22* (11), 1700–1712.

Domínguez-Iino, Tomás, “Efficiency and redistribution in environmental policy: An equilibrium analysis of agricultural supply chains," January 2021.

Donoso, Guillermo, “The evolution of water markets in Chile," in “Water Trading and Global Water Scarcity," Routledge, 2012, pp. 131–149.

Dubois, Olivier et al., *The state of the world's land and water resources for food and agriculture: managing systems at risk.*, Earthscan, 2011.

d'Odorico, Paolo, Joel Carr, Carole Dalin, Jampel Dell'Angelo, Megan Konar, Francesco Laio, Luca Ridolfi, Lorenzo Rosa, Samir Suweis, Stefania Tamea et al., “Global virtual water trade and the hydrological cycle: patterns, drivers, and socio-environmental impacts," *Environmental Research Letters*, 2019, *14* (5), 053001.

Easter, K William and Qiuqiong Huang, “Water markets for the 21st century," VV. AA., *Water Markets: How Do We Expand Their Use*, 2014, p. 1.

_ , Mark W Rosegrant, and Ariel Dinar, *Formal and informal markets for water: institutions, performance, and constraints*, Routledge, 2018.

Eaton, Jonathan and Samuel Kortum, “Technology, geography, and trade," *Econometrica*, 2002, *70* (5), 1741–1779.

Fan, Y., H. Li, and G. Miguez-Macho, “Global patterns of groundwater table depth," *Science*, 2013, *339* (6122), 940–943.

Farrokhi, Farid, Elliot Kang, Heitor Pellegrina, and Sebastian Sotelo, “Deforestation: A Global and Dynamic Perspective," Working Paper, 2023.

Grafton, R Quentin, Gary Libecap, Samuel McGlennon, Clay Landry, and Bob O'Brien, “An integrated assessment of water markets: a cross-country comparison," *Review of Environmental Economics and Policy*, 2020.

Griffin, Ronald C and Gregory W Characklis, “Issues and trends in Texas water marketing," *Journal of Contemporary Water Research and Education*, 2011, *121* (1), 5.

Healy, Stephen, Richard Pearce, and Michael Stockbridge, *The implications of the Uruguay Round Agreement on Agriculture for developing countries: a training manual*, Vol. 41, Food & Agriculture Org., 1998.

Hearne, Robert R and José L Trava, “Water markets in Mexico: opportunities and constraints," Technical Report 1997.

Hendricks, Nathan P, Aaron D Smith, and Nelson B Villoria, \Global Agricultural Supply Response to Persistent Price Shocks," 2018.

_ , Aaron Smith, and Daniel A Sumner, \Crop supply dynamics and the illusion of partial adjustment," American Journal of Agricultural Economics, 2014, 96 (5), 1469{1491.

_ and Jeffrey M Peterson, \Fixed effects estimation of the intensive and extensive margins of irrigation water demand," Journal of Agricultural and Resource Economics, 2012, pp. 1{19.

Hengl, Tomislav, Jorge Mendes de Jesus, Gerard B. M. Heuvelink, Maria Ruiperez Gonzalez, Milan Kilibarda, Aleksandar Blagotic, Wei Shangguan, Marvin N. Wright, Xiaoyuan Geng, Bernhard Bauer-Marschallinger, Mario Antonio Guevara, Rodrigo Vargas, Robert A. MacMillan, Niels H. Batjes, Johan G. B. Leenaars, Eloi Ribeiro, Ichsani Wheeler, Stephan Mantel, and Bas Kempen, \SoilGrids250m: Global gridded soil information based on machine learning," PLOS ONE, February 2017, 12 (2).

Hoekstra, Arjen Y. and Mesfin M. Mekonnen, \The water footprint of humanity," Proceedings of the National Academy of Sciences, 2012, 109 (9), 3232{3237.

Hsiao, Allan, \Coordination and commitment in international climate action: Evidence from palm oil," January 2021. Job market paper (MIT).

Kloezen, Wim H, \Water markets between Mexican water user associations," Water Policy, 1998, 1 (4), 437{455.

Konar, Megan, Zekarias Hussein, Naota Hanasaki, Denise L Mauzerall, and Ignacio Rodriguez-Iturbe, \Virtual water trade flows and savings under climate change," Hydrology and Earth System Sciences, 2013, 17 (8), 3219{3234.

Libecap, Gary D, \Transaction costs, property rights, and the tools of the new institutional economics: water rights and water markets," New institutional economics: a guidebook, 2008, pp. 272{291.

Loheide, Steven P., James J. Butler, and Steven M. Gorelick, \Estimation of groundwater consumption by phreatophytes using diurnal water table fluctuations: A saturated-unsaturated flow assessment," Water Resources Research, July 2005, 41 (7).

Mas-Colell, Andreu, Michael D. Whinston, and Jerry R. Green, Microeconomic theory, Oxford: Oxford University Press, 1995.

Mekonnen, M. M. and A. Y. Hoekstra, “The green, blue and grey water footprint of crops and derived crop products,” Hydrology and Earth System Sciences, May 2011, 15 (5), 1577–1600.

Monfreda, Chad, Navin Ramankutty, and Jonathan A. Foley, “Farming the planet: Geographic distribution of crop areas, yields, physiological types, and net primary production in the year 2000,” Global Biogeochemical Cycles, March 2008, 22 (1).

Palomo-Hierro, Sara, José A Gómez-Limón, and Laura Riesgo, “Water markets in Spain: Performance and challenges,” Water, 2015, 7 (2), 652–678.

Pekel, Jean-François, Andrew Cottam, Noel Gorelick, and Alan S. Belward, “High-resolution mapping of global surface water and its long-term changes,” Nature, 2016, 540, 418–422.

Phillips, Keith R, Judy Teng et al., “Groundwater Markets Slowly Evolve in Ever-Thirstier Texas,” Southwest Economy, 2020, (First Quarter).

Portmann, Felix T., Stefan Siebert, and Petra Döll, “MIRCA2000—Global monthly irrigated and rainfed crop areas around the year 2000: A new high-resolution data set for agricultural and hydrological modeling,” Global Biogeochemical Cycles, March 2010, 24 (1).

Potapov, Peter, Svetlana Turubanova, Matthew C Hansen, Alexandra Tyukavina, Viviana Zalles, Ahmad Khan, Xiao-Peng Song, Amy Pickens, Quan Shen, and Jocelyn Cortez, “Global maps of cropland extent and change show accelerated cropland expansion in the twenty-first century,” Nature Food, 2022, 3 (1), 19–28.

Rafey, Will, “Droughts, deluges, and (river) diversions: Valuing market-based water reallocation,” American Economic Review, 2023, 113 (2), 430–471.

Richey, Alexandra S, Brian F Thomas, Min-Hui Lo, John T Reager, James S Famiglietti, Katalyn Voss, Sean Swenson, and Matthew Rodell, “Quantifying renewable groundwater stress with GRACE,” Water resources research, 2015, 51 (7), 5217–5238.

Roberts, Michael J and Wolfram Schlenker, “Identifying supply and demand elasticities of agricultural commodities: Implications for the US ethanol mandate,” American Economic Review, 2013, 103 (6), 2265–95.

Rodell, Matthew, Jay S Famiglietti, David N Wiese, JT Reager, Hiroko K Beaudoing, Felix W Landerer, and M-H Lo, “Emerging trends in global freshwater availability,” Nature, 2018, 557 (7707), 651–659.

Rossi-Hansberg, Esteban, \Geography of growth and development," in \Oxford Research Encyclopedia of Economics and Finance," Oxford University Press, July 2019.

Ryan, Nicholas and Anant Sudarshan, \Rationing the commons," Journal of Political Economy, November 2021, forthcoming.

Saleth, Rathinasamy Maria, Strategic analysis of water institutions in India: Application of a new research paradigm, Vol. 79, IWMI, 2004.

Schwabe, Kurt, Mehdi Nemati, Clay Landry, and Grant Zimmerman, \Water markets in the Western United States: Trends and opportunities," Water, 2020, 12 (1), 233.

Scott, Paul, \Dynamic discrete choice estimation of agricultural land use," 2014.

Sekhri, Sheetal, \Wells, water, and welfare: The impact of access to groundwater on rural poverty and conflict," American Economic Journal: Applied Economics, July 2014, 6 (3), 76{102.

_, \Agricultural trade and depletion of groundwater," Journal of Development Economics, 2022, 156, 102800.

Sheffield, Justin, Gopi Goteti, and Eric F Wood, \Development of a 50-year high-resolution global dataset of meteorological forcings for land surface modeling," Journal of climate, 2006, 19 (13), 3088{3111.

Tapley, Byron D., Srinivas Bettadpur, John C. Ries, Paul F. Thompson, and Michael M. Watkins, \GRACE measurements of mass variability in the earth system," Science, July 2004, 305 (5683), 503{505.

Taylor, Charles A, \Irrigation and Climate Change: Long-Run Adaptation and its Externalities," NBER Chapters, 2022.

Trabucco, Antonio and Robert Zomer, \Global Aridity Index and Potential Evapotranspiration (ETO) Climate Database v2," 2019.

Wang, Xin and Shizhong Yang, \Research on status quo of water rights trading in China," in \IOP Conference Series: Earth and Environmental Science," Vol. 171 IOP Publishing 2018, p. 012033.

Young, Mike, \Trading into and out of trouble: Australian's water allocation and trading experience," in \Water Trading and Global Water Scarcity," Routledge, 2012, pp. 114{130.

Zhang, Lijuan, Jinxia Wang, Jikun Huang, and Scott Rozelle, \Development of groundwater markets in China: a glimpse into progress to date," World Development, 2008, 36 (4), 706{726.