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

Tamma Carleton† UC Berkeley and NBER
Levi Crews‡ UCLA
Ishan Nath§ San Francisco Fed

20 July 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

*We thank Jonathan Dingel, José-Antonio Espín-Sánchez, Farid Farrokhi, Scott Jasechko, Sam Kortum, Andrés Rodríguez-Clare, Joe Shapiro, Felix Tintelnot, and seminar participants at CalTech, FRB St. Louis, Georgetown, Harvard, Johns Hopkins, Michigan, Minnesota, MIT, NYU, Penn State, Princeton, UBC, UC Berkeley, UC Davis, UCLA, UChicago, USC, the World Bank, Yale, the BFI Coase Project, the Fed Junior Applied Group, the LACEA Annual Meeting, LSE Environment Week, the SED, and the Virtual Workshop on Trade, Spatial Economics, and the Environment for helpful comments. Crews and Nath thank the International Economics Section at Princeton University for their hospitality while working on this project. Zoë Arnaut-Hull, Cheikh Fall, Yang Gao, Ethan Goode, Christian Hilgemann, and Walker Lewis provided excellent research assistance. All remaining errors are our own. 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.

†tcarleton@ucsb.edu
‡lgcrews@econ.ucla.edu
§inath33@gmail.com

1 Introduction

Over 90% of global water use by humans occurs in agricultural production, which is critically dependent on rainfall and local stocks of ground 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 in output markets 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 agri-

cultural policy. New datasets originating from hydrology, agronomy, and remote sensing indicate the availability of ground and surface water stocks, trends in water availability over time, and the spatial distribution of agricultural production, irrigation, water consumption, and productivity. Administrative records capture prices, trade patterns, and the wide array of policy interventions shaping agricultural markets. To our knowledge, this dataset constitutes the largest collection of global data on water and agriculture to be used in economics.

We use these data and corresponding scientific literature to establish a set of five facts about water, agriculture, and trade that frame the subsequent 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. Notably, this finding that specialization follows factor abundance occurs despite the well-documented distortions both over water as an input and over the agricultural output markets shaping global trade patterns.

To fully characterize the role of trade in governing global water resources and explore policy counterfactuals, we build the first dynamic spatial equilibrium model

of agricultural production, consumption, and trade that incorporates water resources and includes 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 our 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 90$^{\text{th}}$ 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, 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

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

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 private 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*, partially, reflect its scarcity. In this setting, while spatial and temporal externalities abound, the natural environment partially 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, the predominance of input and output distortions implies that specific agricultural policy liberalizations can have the opposite effect and exacerbate depletion and lower welfare. 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 uses a wide array of geospatial data to establish a set of stylized facts that frame the analysis. Section 3 lays out the model, and Section 4 shows how we calibrate it to match the data. Section 5 shows results from the policy counterfactuals, and Section 6 concludes.

2 Stylized facts

2.1 Five facts about water and agriculture

We assemble a wide array of geospatial datasets to create what constitutes, to our knowledge, the largest collection of global data on water and agriculture yet to be used in economics. This compilation of 18 distinct datasets includes data on water resources and on agricultural land use, production, productivity, and policy and is composed of over 5 billion observations. It is detailed in Appendix A. Here, we bring this data together 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 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 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. Using standard gridded climate data from Sheffield, Goteti and Wood (2006), we estimate that the median parcel of global arable land receives 7,021 cubic meters per hectare (m^3/ha) of rainfall per year, and the 1^{st} to 99^{th} percentile ranges from 44 to 33,208 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, while the least water-intensive staple, millet, requires just 4,300 m^3/ha.[2] Less than 42% of the world's arable land receives enough rainfall to supply full coverage of rice acreage without drawing down the stock of water in steady state. In contrast, the driest 31% of global arable land receives less rainfall than is required even to cover all acreage with millet.[3] To be sure, these proportions do not represent a binding limit on the share of regions that can grow each of these crops

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

[3]These and related statements are highly conservative estimates due to the presence of water runoff. 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 rates exist, but accounting for runoff would imply that fewer regions of the globe have sufficient rainfall to supply the most water-intensive 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, as revealed by high-resolution gridded estimates of groundwater table depth assembled by combining a hydrologic model with over 1.6 million well observations (Fan et al., 2013). In addition to the 15% of the world's surface area that are covered by local bodies of surface water (e.g. small lakes, rivers, and wetlands) that can be used for irrigation, 8% of arable land area has easily accessible groundwater within one meter of the earth's surface. On the other hand, 64% 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 these data identify only the top of the water table, this information is what determines the costs of extracting water for farmers, as we detail when developing the quantitative model in Section 3.

A dynamic measure of trends in total water storage also varies widely across the world's arable land, previewing future availability if current patterns continue. We obtain estimates of changes in total water storage from the Gravity Recovery and Climate Experiment (GRACE). This satellite-based measure quantifies changes over time in all water in a region, including both groundwater and surface water, at a ∼1 degree resolution and with monthly frequency (see Appendix A for details). In total, over the satellite record of 2003–2016, just over 50% of the world's arable land was losing water, while the other half was gaining water. For instance, the southwestern U.S., the Middle East, and northern India are all losing water, while western Europe, eastern Australia, and southwestern Africa are all gaining. It is worth noting that the magnitudes of these trends in total water storage are rather modest. The 1^{st} to 99^{th} percentile of arable land ranges from losing 224 m^3/ha to gaining 188 m^3/ha. For context, median global rainfall on arable land is about 7,021 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 agriculture.

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 leveraging hydrologic and agronomic modeling in combination with detailed spatial data on climate, soils, and agricultural practices indicates that 92% of global *consumptive* water use by humans is dedicated to agricultural production (Hoekstra and Mekonnen, 2012; d'Odorico et al., 2019).[4] In contrast, industrial production accounts for only 4.4% and domestic water supply for direct commercial and residential use for just 3.6%.

A significant amount of this water consumption is implicitly traded through the global agricultural market. Hoekstra and Mekonnen (2012) and Carr et al. (2013) 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, these findings 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 (Libecap, 2008). 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 use or other institutional arrangements without a clear market-based mechanism for allocating the resource to its highest value uses. Recent research in Australia and the U.S. shows that the implementation of water markets generates large economic benefits by allowing water to flow to those uses with the highest value, reducing the total extraction necessary to achieve a given level of output (Rafey, 2023; Ayres et al., 2021; Bruno and Jessoe, 2021).

Despite the theoretically and empirically documented benefits of tradable property rights for water, the existence of such markets remains exceedingly rare. We

[4]See Appendix A for details on how these estimates are constructed in the agronomic literature.

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 (Appendix Figure C.1). Formally established tradable property rights for water exist in Australia, Chile, Mexico, Spain, South Africa, Oman, and a number of states in the western U.S.[5] 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 nearly all cases). Some additional water trading occurs outside established formal markets in parts of India, Pakistan, and China, though without clear mechanisms to resolve the common pool resource challenge (Saleth, 2004; Zhang et al., 2008; Easter et al., 2018).[6] 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, implementation challenges are ubiquitous. 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).[7] 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, but the widespread implementation of such 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 3 would be to eliminate both input side market failures and all distortions affecting output markets. As long as local property rights market failures remain

[5]U.S. states with formal water markets for farmers include Arizona, California, 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, Libecap, McGlennon, Landry and O'Brien (2020) for broad global overviews, Donoso (2012) for more details on Chile, Hearne and Trava (1997) and Kloezen (1998) on Mexico, Palomo-Hierro, Gómez-Limón and Riesgo (2015) on Spain, and Young (2012) on Australia.

[6]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).

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

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 documents the critical role that crop subsidies, taxes, tariffs, and other price shocks play in driving farmer decisions, such as land use, and farm outcomes, such as supply and productivity (e.g., Roberts and Schlenker, 2013; Hendricks et al., 2014; Scott, 2014; Hendricks et al., 2018). 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 government crop price interventions—measured with an aggregate metric capturing the wedge between distorted and counterfactual undistorted prices called the "nominal rate of assistance" (NRA, detailed in Appendix A)—on ΔTWS 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 25^{th} or 75^{th} percentile. Similarly, Sekhri (2022) finds that Indian policies that promote agricultural trade have large negative consequences for groundwater reserves.

These estimates suggest that even modest changes in agricultural policy can play a critical role in shaping the evolution of water availability over time. Such agricultural market interventions are ubiquitous around the world (Anderson et al., 2013), yet 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: Specialization in water-intensive crops follows water abundance, with important exceptions

If global water use were efficiently allocated, 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 descriptively evaluate this hypothesis by using the spatial datasets described in Appendix A to examine the global relationship between the water intensity of land use and water availability. We start by calculating a simple index of the water intensity of land use across all crops. To do so, we use 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 128 crops that account for over 88% 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. We then show how this index water intensity of arable land correlates spatially with each of the three main measures of water resources described above: rainfall, groundwater table depth, and changes in total water storage (Δ TWS).

Figure 2 shows the results. Maps on the left overlay water intensity of arable land with each of the three water availability measures, while scatter plots on the right show average water intensity of land use by decile of water availability. The results suggest that water-intensive agricultural production strongly clusters in those regions with high water tables and ample rainfall. The maps in panels (a) and (c) indicate that most regions of low water abundance also exhibit very low agricultural water use (regions in grey, such as the western U.S., much of Australia, and global deserts), while regions with ample water resources plant highly water-intensive crops (regions in yellow, such as southeast Asia, the midwestern U.S., and Amazonia). Scatterplots in panels (b) and (d) show that 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 three times as water-intensive as that of the lowest quintile. Even more strikingly, the rainiest quintile of places in the world use nine 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 25% and 100% 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 2, panels (e) and (f). In contrast to the static measures of groundwater

and rainfall in which water-intensive production is highly concentrated in water-abundant regions, panels (e) and (f) show that the regions in the world *losing* water most rapidly in the period from 2003–2016 (such as northern India, Turkey, and northeastern China) have the *highest* water intensity of production. While the average acre of arable land in the world uses about 2400 m^3/ha of water, the average acre in the bottom decile of ΔTWS uses over 3000 m^3/ha. Even more dramatically, the average acre in the bottom 1% of the ΔTWS distribution uses nearly 3500 m^3/ha of water, almost 42% higher than the world average. This result aggregates over highly heterogeneous patterns across water intensive crops. For example, rice is disproportionately planted in the world's wettest regions and is exceedingly rare in the world's driest regions (Appendix Figure C.2(a)). Conversely, almonds, among the most water-intensive specialty crops, are primarily grown in dry places, consistent with the aggregate patterns shown in Figure 2. Thus, while the first four panels of Figure 2 show that water-intensive production is concentrated in regions that are currently water-abundant, the last two suggest that some isolated locations planting water-intensive crops could be at risk of future depletion if current trends continue. However, the concurrence of water-intensive production and rapid water loss is restricted to a small number of regions, highlighted in bright red on the map in panel (e).

To take a particularly extreme example, the Indian regions in the bottom 1% of ΔTWS, which are primarily located near the capital of New Delhi and in the northeastern state of Uttarakhand, have crops planted on over 91% of their arable land with an average water intensity of land use of over 8700 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. We explore the characteristics of the places losing water rapidly in much more detail in Carleton et al. (2024).

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 3 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 ΔTWS, and that regions with low rainfall and low or declining stocks have very low water intensity, consistent with the overall results in Figure 2. However, this figure also shows 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 6.8% of the world's arable land.

Overall, we uncover a dominant pattern of water-intensive global agricultural production concentrating in presently water-abundant regions, although a small

share of regions 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.

3 Model

3.1 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.[8] 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.[9] Fields comprise a continuum of heterogeneous parcels, indexed by ω. 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.[10] A farmer uses his own labor to extract groundwater, which he combines with his land and remaining time to produce one of multiple crops, indexed by $k \in \mathcal{K} \equiv \{1, \ldots, K\}$.

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

[9] 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.

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

3.1.1 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_{it}^o, and an agricultural composite, C_{it}:

$$U_{it} = C_{it}^o + \zeta_i \ln C_{it}. \tag{1}$$

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, $\zeta_i \geq 0$.

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

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

and

$$C_{it}^k = \left[\sum_{j \in \mathcal{I}} \left(\zeta_{ji}^k\right)^{1/\sigma} \left(C_{jit}^k\right)^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}}, \tag{3}$$

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

3.1.2 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 $L_t^{fk}(\omega)$ with groundwater $G_t^{fk}(\omega)$ and a fraction $H_t^{fk}(\omega)$ of his labor endowment, the farmer of parcel ω can produce

$$Q_t^{fk}(\omega) = A^{fk}(\omega) \left[H_t^{fk}(\omega)\right]^{\alpha} \times \left[\min\left\{L_t^{fk}(\omega), \frac{G_t^{fk}(\omega)}{\phi^k}\right\}\right]^{1-\alpha}, \tag{4}$$

where $A^{fk}(\omega) \geq 0$ denotes the total factor productivity of parcel ω in field f if allocated to crop k, and ϕ^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.[11]

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:

$$\mathbb{P}\left\{A^{f1}(\omega) \leq a^1, \ldots, A^{fK}(\omega) \leq a^K\right\} = \exp\left\{-\gamma \sum_{k \in \mathcal{K}} \left(\frac{a^k}{A^{fk}}\right)^{-\theta}\right\} \tag{5}$$

where $\theta > 1$ measures the extent of technological heterogeneity within each field and the constant γ is set such that $A^{fk} = \mathbb{E}[A^{fk}(\omega)]$.[12] The term $A^{fk} \geq 0$ 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 $1 - H_t^{fk}(\omega)$ 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

$$A_q^w(D_{qt}) = \Upsilon_q D_{qt}^{-\upsilon}. \tag{6}$$

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 Υ_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 θ:

$$\mathbb{P}\left\{A_i^o(\omega) \leq a^o\right\} = \exp\left\{-\gamma \left(\frac{a^o}{A_i^o}\right)^{-\theta}\right\} \tag{7}$$

where $A_i^o = \mathbb{E}[A_i^o(\omega)]$ is the average labor productivity in country i's outside sector.[13] Importantly, draws from this distribution are independent of the crop-specific productivity draws for the parcel the worker would otherwise till.

[11] 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 B.1.5 discusses a nested CES technology that generalizes that in (4).

[12] Formally, we set $\gamma \equiv \Gamma[(\theta - 1)/\theta]^{-\theta}$, where $\Gamma(\cdot)$ denotes the gamma function.

[13] Some readers may find it odd that workers have heterogeneous productivities in the outside

3.1.3 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 δ_{ij}^k 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

$$p_{ijt}^k = \delta_{ij}^k p_{it}^k, \tag{8}$$

where p_{it}^k 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 τ_{it}^k for each crop k at each date t.[14] Accordingly, the distorted farm gate price that farmers receive is $\tau_{it}^k p_{it}^k$, and the distorted price paid by the consumer in j is $\tau_{it}^k p_{ijt}^k = \delta_{ij}^k \tau_{it}^k p_{it}^k$. 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

$$D_{qt+1} = D_{qt} + \rho_q[(1-\psi)X_{qt} - R_q] \tag{9}$$

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), ψ is the constant rate of return flow,[15] and ρ_q is an aquifer-specific conversion factor between volume and depth that depends on the local soil type.

3.2 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.

---

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.

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

[15] 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 ψ accounts for the latter fraction.

3.2.1 Utility maximization

Given equations (1), (2), (3), and (8), utility maximization by the representative household in each country requires that[16]

$$C_{jit}^k = \zeta_i \frac{\zeta_i^k \left(P_{it}^k\right)^{1-\kappa}}{\sum_{\ell \in \mathcal{K}} \zeta_i^\ell \left(P_{it}^\ell\right)^{1-\kappa}} \frac{\zeta_{ji}^k \left(\delta_{ji}^k \tau_{jt}^k p_{jt}^k\right)^{-\sigma}}{\sum_{n \in \mathcal{I}} \zeta_{ni}^k \left(\delta_{ni}^k \tau_{nt}^k p_{nt}^k\right)^{1-\sigma}} \quad \text{for all } i, j \in \mathcal{I}, \;\; k \in \mathcal{K}, \quad (10)$$

where

$$P_{it}^k = \left[\sum_{n \in \mathcal{I}} \zeta_{ni}^k \left(\delta_{ni}^k \tau_{nt}^k p_{nt}^k\right)^{1-\sigma}\right]^{\frac{1}{1-\sigma}}$$

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

3.2.2 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 $L_t^{fk}(\omega) = 1$ 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 $A^{fk}(\omega) M(\phi^k, D_{q(f)t})$, where M is decreasing in the water intensity ϕ^k and the current water table depth $D_{q(f)t}$.[17] Thus, the farmer's revenue from growing crop ℓ is

$$r_t^{f\ell}(\omega) = \tau_{i(f)t}^\ell p_{i(f)t}^\ell A^{f\ell}(\omega) M(\phi^k, D_{q(f)t}).$$

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

$$\pi_t^{fk} \equiv \mathbb{P}\left\{r_t^{fk}(\omega) = \max\{A_{i(f)}^o(\omega), r_t^{f1}(\omega), \ldots, r_t^{fK}(\omega)\}\right\}.$$

[16] See B.1.1 for derivation.

[17] See the definition and derivation of M in B.1.2. An important point to note is that $\lim_{D \to \infty} M(\phi^k, D) = 0$ 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 $t \in \mathbb{N}$,[18]

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

where

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

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, ϕ^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.[19] 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

$$Q_{it}^k = \sum_{f \in \mathcal{F}_i} h^f \pi_t^{fk} M(\phi^k, D_{q(f)t}) \mathbb{E}\left[A^{fk}(\omega) \middle| r_t^{fk}(\omega) = \max\{A_{i(f)}^o(\omega), r_t^{f1}(\omega), \ldots, r_t^{fK}(\omega)\}\right].$$

[18] See B.1.3 for derivation.

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

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

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[20]

$$\mathbb{E}\left[A^{fk}(\omega)\middle| r_t^{fk}(\omega) = \max\{A^o_{i(f)}(\omega), r_t^{f1}(\omega), \ldots, r_t^{fK}(\omega)\}\right] = A^{fk}\left(\pi_t^{fk}\right)^{-1/\theta}.$$

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:

$$Q_{it}^k = \sum_{f \in \mathcal{F}_i} h^f A^{fk} M(\phi^k, D_{q(f)t}) \left(\pi_t^{fk}\right)^{\frac{\theta-1}{\theta}} \tag{12}$$

for all $i \in \mathcal{I}$ and $k \in \mathcal{K}$.

3.2.3 Market clearing and feasibility

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

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

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

$$X_{qt} = \sum_{f \in \mathcal{F}_q} \sum_{k \in \mathcal{K}} h^f x_t^{fk} \pi_t^{fk}. \tag{14}$$

where x_t^{fk} is optimal water extraction to grow crop k on field f.[21]

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.

3.2.4 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, $\{\tau_{it}^k\}$, 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, $\{\pi_t^{fk}\}$, groundwater depths,

[20] See B.1.4 for derivation.

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

$\{D_{qt}\}$, and groundwater extractions, $\{X_{qt}\}$, such that equations (9), (10), (11), (12), (13), and (14) hold.[22]

Conditions for the existence and uniqueness of an equilibrium are established in Appendix B.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.[23] 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.[24]

4 Estimation

To simulate the model described in Section 3, 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 describe the data we use and the sample restrictions that we impose on the inputs to our estimation procedure. Details on the datasets used can be found in Appendix A. We conclude by checking the model's fit of targeted and untargeted moments.

4.1 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 (which measure changes in water storage over time) 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.

[22] 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 B.2.4 for discussion.

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

[24] 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.

To construct our sample, we select 52 countries ($i \in \mathcal{I}$) that account for 94% of total GDP, 97% of total agricultural production value, and 99% of total agricultural laborers (see Figure 4a).[25] Within these countries, we model land use at the resolution of 5 arc-minute grid cells to match the FAO's Global Agro-Ecological Zones (GAEZ) dataset, which contains potential yields for 38 crops using a spatially explicit agronomic model. We restrict our attention to arable land, which leaves us with roughly 1.9 million "fields". From the set of crops in GAEZ, we select 22 crops ($k \in \mathcal{K}$) that account for 56% of total agricultural production and 59% of total water use. Our set of crops, shown in Figure 4c, 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" ($q \in \mathcal{Q}$) as clusters of contiguous GRACE grid cells 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 grid cells 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 grid cells 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 grid cells that lie fully outside the borders of our sample countries. The surviving 205 clusters map to what we call "aquifers" in the model.

4.2 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 Equation (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 σ between different varieties of a given crop and to invert a composite of the crop-origin-specific demand shifters, ζ_{ji}^k, and trade costs, δ_{ji}^k.[26] Second, we

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

[26] 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

use the previous estimates to construct price indices at the crop level, P_i^k, and combine them with data on crop expenditures, $E_i^k = \sum_{j \in \mathcal{I}} E_{ji}^k$, to estimate the elasticity of substitution κ between crops and to invert the crop-specific demand shifters, ζ_i^k. Finally, we use data on total crop expenditures, $E_i = \sum_{k \in \mathcal{K}} E_i^k$, to invert the uppermost demand shifters, ζ_i. Throughout, we will make repeated use of the identity

$$E_{ji}^k = \left(\delta_{ji}^k p_j^k\right) C_{ji}^k = \zeta_i \frac{\zeta_i^k \left(P_i^k\right)^{1-\kappa}}{\sum_{\ell \in \mathcal{K}} \zeta_i^\ell \left(P_i^\ell\right)^{1-\kappa}} \frac{\zeta_{ji}^k \left(\delta_{ji}^k p_j^k\right)^{1-\sigma}}{\sum_{n \in \mathcal{I}} \zeta_{ni}^k \left(\delta_{ni}^k p_n^k\right)^{1-\sigma}}, \tag{15}$$

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 $\zeta_{ji}^k (\delta_{ji}^k)^{1-\sigma} = 0$. 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

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

will be captured by an importer-crop fixed effect, and the last term $\varepsilon_{ji}^k \equiv \ln[\zeta_{ji}^k (\delta_{ji}^k)^{1-\sigma}]$ 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

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

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,

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

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.[27]

country i's total production and total exports of crop k.

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

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 $\zeta_{ji}^k(\delta_{ji}^k)^{1-\sigma}$ 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.

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,

$$P_i^k = \left[\sum_{j\in\mathcal{I}} \zeta_{ji}^k(\delta_{ji}^k p_j^k)^{1-\sigma}\right]^{\frac{1}{1-\sigma}},$$

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

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

where the first term

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

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

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

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 ζ_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 Equation (1) implies that $\zeta_i = E_i$ for all $i \in \mathcal{I}$. 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.

4.3 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.[28] 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 in Equation (4) to estimates from the literature. First, we set $1-\alpha = 0.25$ 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, ϕ^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. Third, we set the elasticity of labor productivity with respect to depth in the pumping technology (6) to $\upsilon = 1$ 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, θ, each country's average labor productivity in the outside sector, $\mathbf{A^o} \equiv \{A_i^o\}$, and each aquifer's scale of labor productivity in pumping groundwater, $\boldsymbol{\Upsilon} \equiv \{\Upsilon_q\}$. 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 Monfreda, Ramankutty and Foley (2008).

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

$$Q_i^k(\theta, \mathbf{A^o}, \boldsymbol{\Upsilon}) = \sum_{f \in \mathcal{F}_i} h^f A^{fk} M(\phi^k, D_{q(f)}) \left(\frac{\left(\tau_i^k p_i^k A^{fk} M(\phi^k, D_{q(f)})\right)^\theta}{(A_i^o)^\theta + \sum_{\ell \in \mathcal{K}} \left(\tau_i^\ell p_i^\ell A^{f\ell} M(\phi^\ell, D_{q(f)})\right)^\theta} \right)^{\frac{\theta-1}{\theta}}$$

where the dependence on $\boldsymbol{\Upsilon}$ is through the optimal input bundle, M. Likewise, we define the predicted land allocated to agriculture in country i and the predicted

[28] 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.

water extracted from aquifer q as functions of the same parameters in a base year,

$$L_i(\theta, \mathbf{A^o}, \boldsymbol{\Upsilon}) = \sum_{k \in \mathcal{K}} \sum_{f \in \mathcal{F}_i} h^f \frac{\left(\tau_i^k p_i^k A^{fk} M(\phi^k, D_{q(f)})\right)^\theta}{(A_i^o)^\theta + \sum_{\ell \in \mathcal{K}} \left(\tau_i^\ell p_i^\ell A^{f\ell} M(\phi^\ell, D_{q(f)})\right)^\theta}$$

$$X_q(\theta, \mathbf{A^o}, \boldsymbol{\Upsilon}) = \sum_{k \in \mathcal{K}} \sum_{f \in \mathcal{F}_q} h^f x^{fk} \frac{\left(\tau_i^k p_i^k A^{fk} M(\phi^k, D_{q(f)})\right)^\theta}{(A_i^o)^\theta + \sum_{\ell \in \mathcal{K}} \left(\tau_{it}^\ell p_{it}^\ell A^{f\ell} M(\phi^\ell, D_{q(f)t})\right)^\theta}.$$

Note that all three functions encode an equilibrium for a given set of agricultural policies, τ_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.[29]

Conditional on θ, we will search for the vectors $\mathbf{A^o}$ and $\boldsymbol{\Upsilon}$ at which, simultaneously, (i) the total amount of land allocated to crops predicted by the model for each country i, $L_i(\theta, \mathbf{A^o}, \boldsymbol{\Upsilon})$, 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, $X_q(\theta, \mathbf{A^o}, \boldsymbol{\Upsilon})$, exactly matches the total amount of water extraction implied by observed cropped area fractions from Monfreda, Ramankutty and Foley (2008), $\hat{\pi}^{fk}$, and the calibrated water intensity of each crop,

$$\hat{X}_q = \sum_{k \in \mathcal{K}} \sum_{f \in \mathcal{F}_q} h^f \phi^k \hat{\pi}^{fk},$$

expressed in cubic meters. Given $\mathbf{A^o}$ and $\boldsymbol{\Upsilon}$, we then search for the value of θ at which output predicted by the model, $Q_i^k(\theta, \mathbf{A^o}, \boldsymbol{\Upsilon})$, 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 $(\theta, \mathbf{A^o}, \boldsymbol{\Upsilon})$ as the solution of

$$\min_{\theta, \mathbf{A^o}, \boldsymbol{\Upsilon}} \sum_{i \in \mathcal{I}} \sum_{k \in \mathcal{K}} \left[\ln Q_i^k(\theta, \mathbf{A^o}, \boldsymbol{\Upsilon}) - \ln \hat{Q}_i^k\right]^2 \tag{20}$$

subject to

$$L_i(\theta, \mathbf{A^o}, \boldsymbol{\Upsilon}) = \hat{L}_i \quad \forall i \in \mathcal{I}, \tag{21}$$

$$X_q(\theta, \mathbf{A^o}, \boldsymbol{\Upsilon}) = \hat{X}_q \quad \forall q \in \mathcal{Q}. \tag{22}$$

In the current draft, we consider only a restricted version of the NLS procedure just described. For now we simply calibrate $\theta = 2.46$ to match the analogous estimate from Costinot, Donaldson and Smith (2016).[30] We then search for the vectors

[29] Nominal rates of assistance are reported in percentage terms, so $\tau_i^k = 1 + \text{NRA}_i^k$. NRAs are not reported for every country-crop pair in 2009. Where no NRA was reported, we set τ_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 $\tau_i^k = 1$. See Appendix A for details.

[30] 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 $[2.28, 2.62]$. Sotelo (2020), using Peruvian data on crop output disaggregated at the district level, estimates $\theta = 2.06$.

$\mathbf{A^o}$ and $\boldsymbol{\Upsilon}$ that simultaneously minimize the mean squared deviations between the model predictions and the data for land use and water extraction.[31]

4.4 Hydrology

We close the quantitative model by specifying the remaining hydrologic parameters that govern the law of motion in Equation (9). We assume a common return flow rate across crops and aquifers, $\psi = 0.25$, which is the central value reported by Dewandel et al. (2008).[32] 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 ρ_q, which accounts for two characteristics of the aquifer. First, and most obviously, its *area*, which can be computed easily from the underlying GRACE grid cells and expressed in squared meters. Second, its *specific yield*, which is the space available in the soil for water mass to be gained or lost . Specific yield depends on the porosity of the soil type, which itself depends on location and depth.[33] 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 ρ_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.

4.5 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 5a and 5b 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. We are attempting to match all agricultural water use for 115 crops contained in the Monfreda, Ramankutty and

[31] 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.

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

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

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.

The geographic pattern of the simulation errors shown in Figures 5a and 5b 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 5c. 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 5d, 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 Υ 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 Υ values to be correlated with plausibly relevant observable covariates.

Figure 6 shows raw correlations of Υ 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 Υ. 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 Υ and our measure of local surface water presence, and the top middle panel shows a U-shaped relationship between temperature and Υ, which is again consistent with relevant estimates in agronomy and agricultural economics.

The covariates shown in Figure 6 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 Υ, 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, $A_q^w(D_{qt})$, 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 6. 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 4.4 that we infer Υ to match the data on the water-intensity of land use. Recall also the patterns shown in Figure 2 and Fact 5 of Section 2 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.

5 Counterfactuals

5.1 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, δ_{ij}^k, 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 7 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 7 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 7 shows the evolution of average global aquifer depth in the baseline and in autarky. As discussed in Section 2, 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 7c 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 Rodríguez-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 7a. 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 7 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 8a 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 8b 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 2, 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 9 shows changes in water table depth across the world in the baseline and in autarky. Recall from Section 4 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, ρ, 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 10 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 10 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 5). 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 8a and 9b 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.

5.2 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; τ_i^k from Section 3) 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 11a and 11b 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 11a). 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 11b).

The maps in Figure 11c and 11d 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 11d 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.

6 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 Andrés Rodríguez-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.

_ , **Levi Crews, and Ishan Nath**, "Is the world running out of fresh water?," in "AEA Papers and Proceedings," Vol. 114 American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203 2024, pp. 31–35.

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.

Chen, JL, CR Wilson, and BD Tapley, "Contribution of ice sheet and mountain glacier melt to recent sea level rise," *Nature Geoscience*, 2013, *6* (7), 549–552.

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 Blagotić, 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.

Sotelo, Sebastian, "Domestic trade frictions and agriculture," *Journal of Political Economy*, 2020, *128* (7), 2690–2738.

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.

Wahr, John M, Steven R Jayne, and Frank O Bryan, "A method of inferring changes in deep ocean currents from satellite measurements of time-variable gravity," *Journal of Geophysical Research: Oceans*, 2002, *107* (C12), 11–1.

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.

Tables & Figures

Figure 1: Summary of Global Data on Water Resources

(a) Average Rainfall (cm/year)

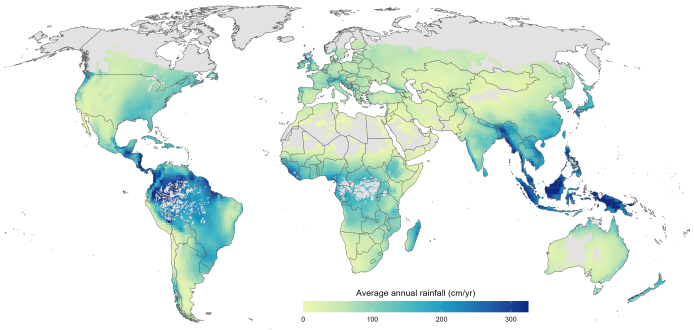

(b) Average Groundwater Table Depth (meters below land surface)

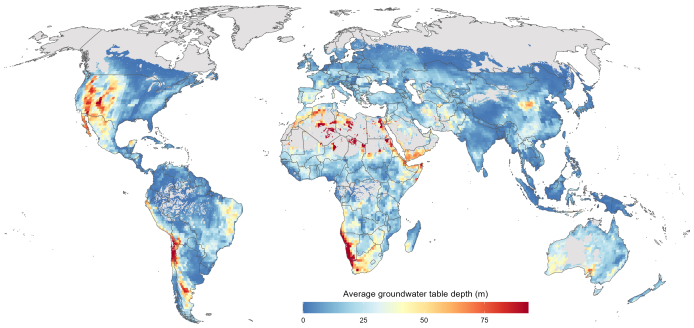

(c) Change in Total Water Storage (cm/year)

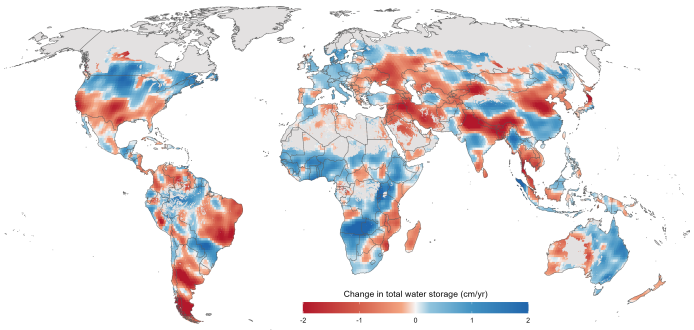

Notes: The map in Panel (a) shows average precipitation at 0.25°×0.25°resolution using data from Sheffield, Goteti and Wood (2006). Panel (b) maps average groundwater measures from Fan, Li and Miguez-Macho (2013) at 30 arc-second resolution. Month-by-month changes in total water storage for equal-area parcels of Earth from the Gravity Recovery and Climate Experiment are shown in Panel (c). See Section A for more details on the water variable data sources.

Figure 2: Correlations of Water Intensity of Arable Land Use and Relative Water Abundance

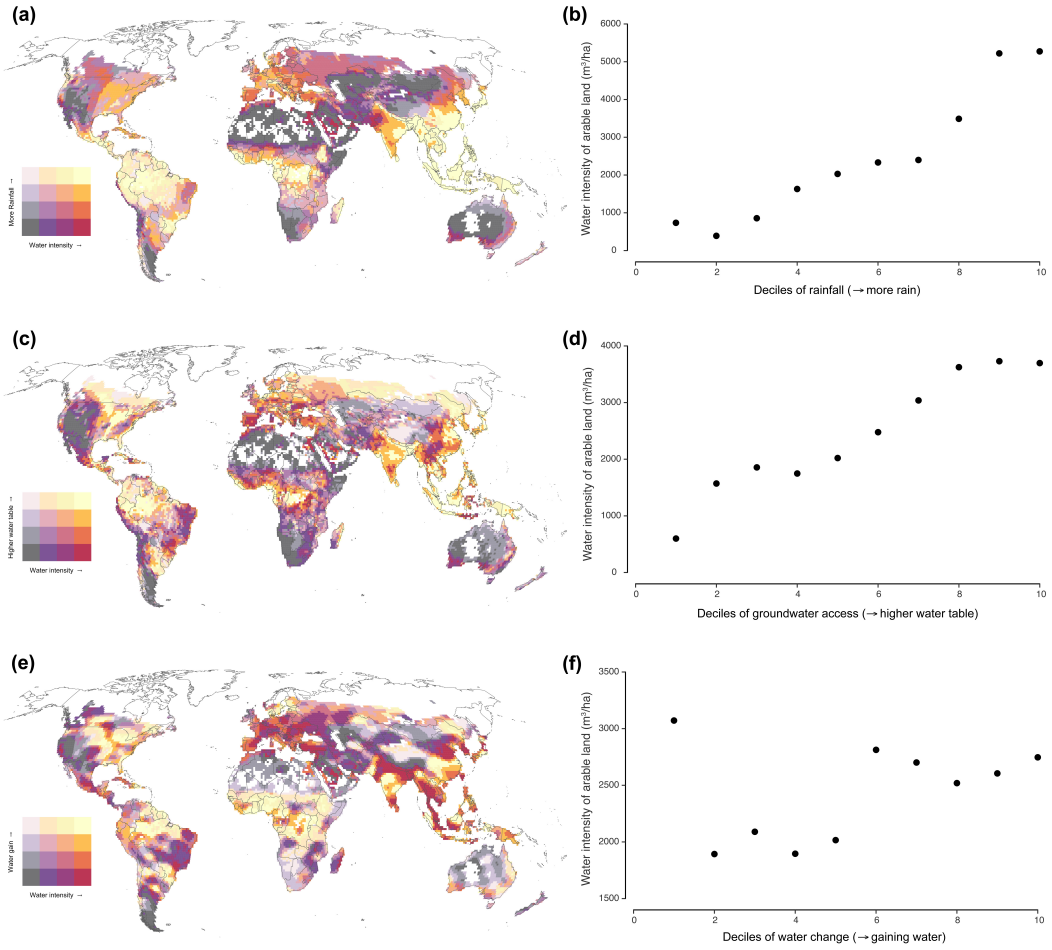

Notes: Maps show water intensity of arable lands against: (a) total annual rainfall from Sheffield et al. (2006); (c) depth to groundwater from Fan et al. (2013); and (e) trends in total water storage from the GRACE satellite. Scatter plots in (b), (d), and (f) show the average water intensity of arable land for each decile of precipitation, groundwater table depth, and change in total water storage, respectively. Arable land is defined as land that is cropped or pastured. The water intensity measures are calculated using data from Mekonnen and Hoekstra (2011), and water variable data sources are detailed in Appendix A.

Figure 3: Water Intensity of Arable Land Use in Water-Stressed Regions

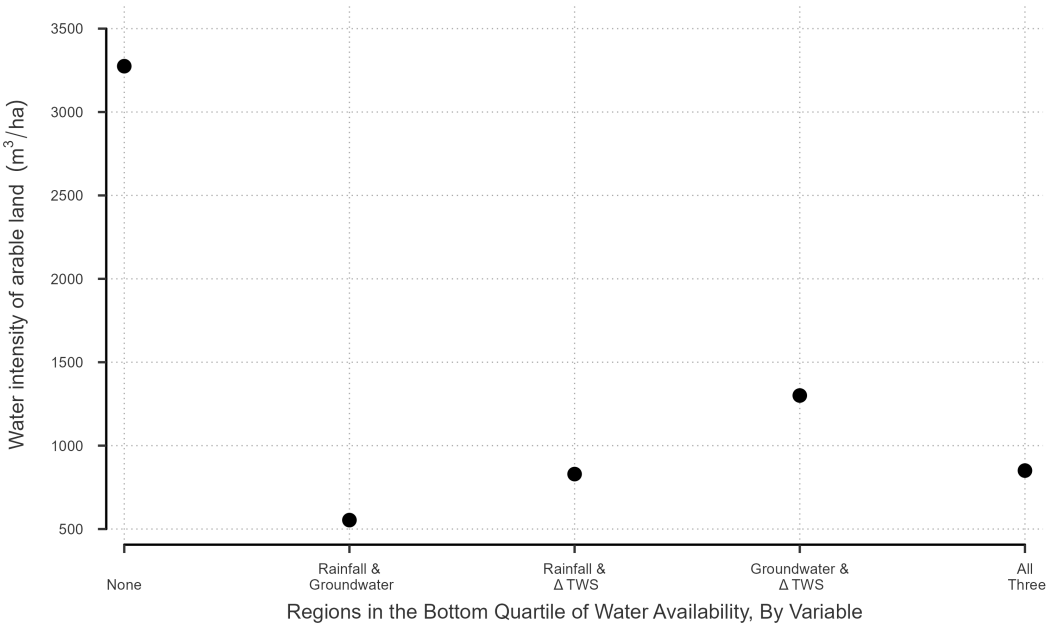

Notes: This graph shows the water intensity of arable land use in regions that fall in the bottom quartile of none, two, or all three of selected water variables: precipitation, groundwater table depth, and change in total water storage. Arable land is defined as land that is cropped or pastured. The water intensity measures are calculated using data from Mekonnen and Hoekstra (2011), and the sources of the water variable data are discussed in Section A.

Figure 4: Selected Sample of Countries, Crops, and Aquifers

(a) 52 countries included in the model

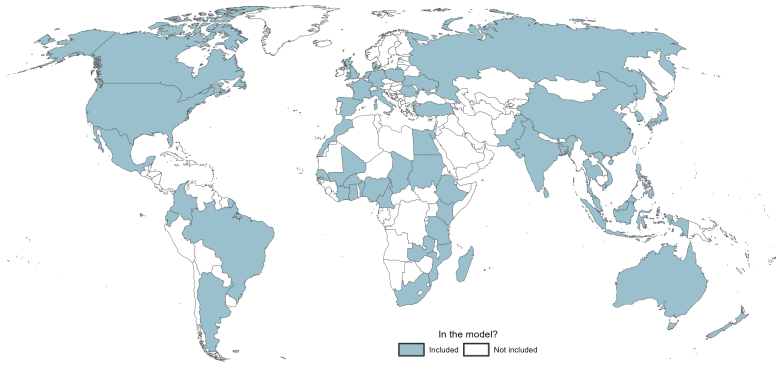

(b) 205 global aquifers included in the model

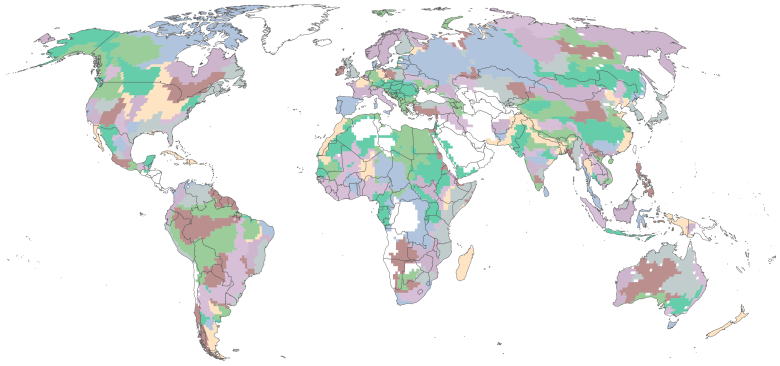

(c) Water intensity of 22 crops included in the model

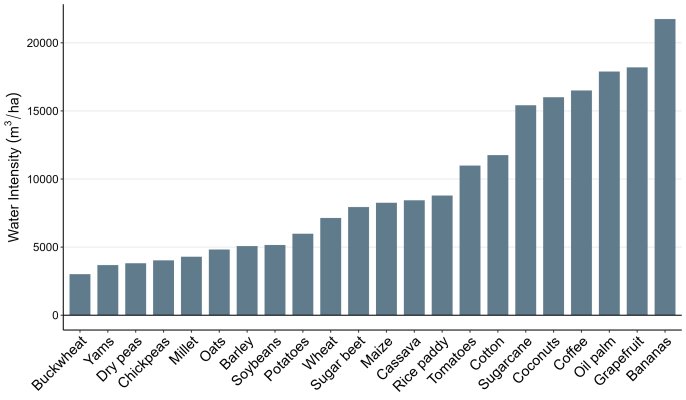

Notes: Panel (a) maps the 52 countries included in the quantitative model. Panel (b) maps 278 global aquifers, 205 of which overlap with the 52 countries shown in panel (a) and are included in the model. Panel (c) shows the water intensity in m^3 per hectare of the 22 crops included in the model. Details on sample selection are provided in Section 4.1.

Figure 5: Simulated Model Fit for Cropped Area and Water Extraction

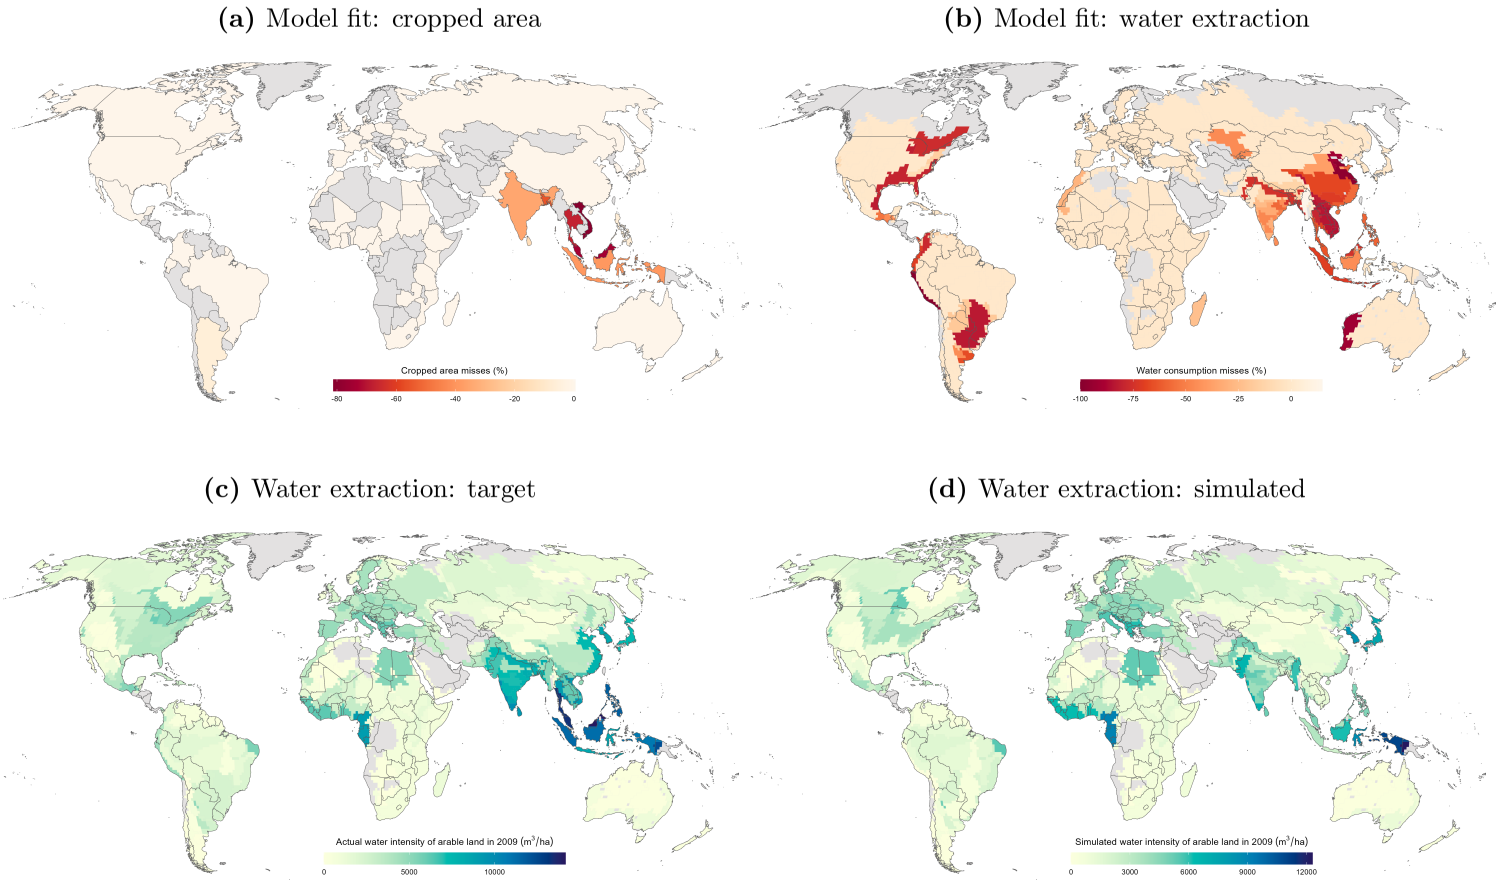

Notes: Panel (a) maps the percentage difference between model-simulated and observed country-level cropped area. Negative values indicate that the model simulation is under-estimating cropped area. Panel (b) maps the percentage difference between model-simulated and observed aquifer-level total water extraction. Negative values indicate that the model simulation is under-estimating water extraction. Panels (c)-(d) map aquifer-level observed (c) and model simulated (d) water intensity of arable land.

Figure 6: Correlations Between Model Parameters Mapping Aquifer Depth to Extraction Productivity (Υ) and Climate and Economic Covariates

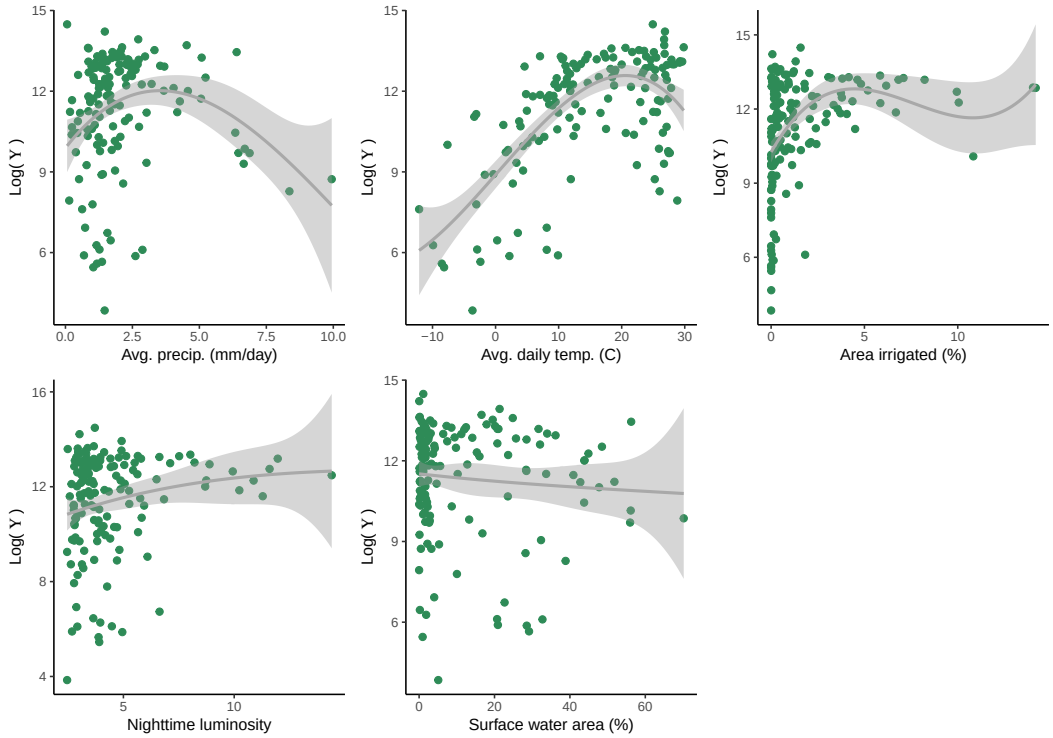

Notes: Each panel in the above figure plots aquifer-level estimates of the logarithm of Υ, the parameter that governs the productivity of water extraction conditional on aquifer depth, against a potential correlate. Variables shown are as follows (left to right, top to bottom): average daily precipitation (mm), average daily temperature (°C), area equipped for irrigation (%), nighttime luminosity (nanowatts per cm^2 per steradian), and surface water area (%).

Table 1: Partial Correlations of Aquifer-Level Covariates, Extraction Productivity ($A_q^w(D_{qt})$), and Impact of Depth on Extraction Productivity (Υ_q)

	Dependent Variable	
	$\log(\Upsilon)$	$\log(A_q^w(D_{qt}))$
Precipitation	0.64**	0.54*
	(0.25)	(0.28)
Precipitation^2	-0.11**	-0.08**
	(0.03)	(0.03)
Temperature	0.26***	0.17***
	(0.04)	(0.05)
Temperature^2	-0.004***	-0.003*
	(0.001)	(0.002)
Area irrigated (%)	0.10*	0.10*
	(0.05)	(0.05)
Nighttime luminosity	0.20***	0.18**
	(0.07)	(0.01)
Surface water area (%)	-0.02**	-0.02*
	(0.01)	(0.01)
Groundwater depth (m)		0.04***
		(0.01)

Notes: Table shows regressions of aquifer-level water extraction productivity, $A_q^w(D_{qt})$, and impact of depth on extraction productivity, Υ, on the same set of covariates as shown in Figure 6. Equation 6 shows the mapping between groundwater depth, Υ, and extraction productivity. The Υ parameter captures heterogeneity across aquifers in the effective input costs of water conditional on water table depth, potentially caused by factors such as rainfall, surface water, irrigation technology, and institutions. Standard errors are shown in parentheses, and stars denote $^*p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$.

Figure 7: Simulated Impact of Agricultural Trade on Global Agriculture, Water, and Welfare Over Time

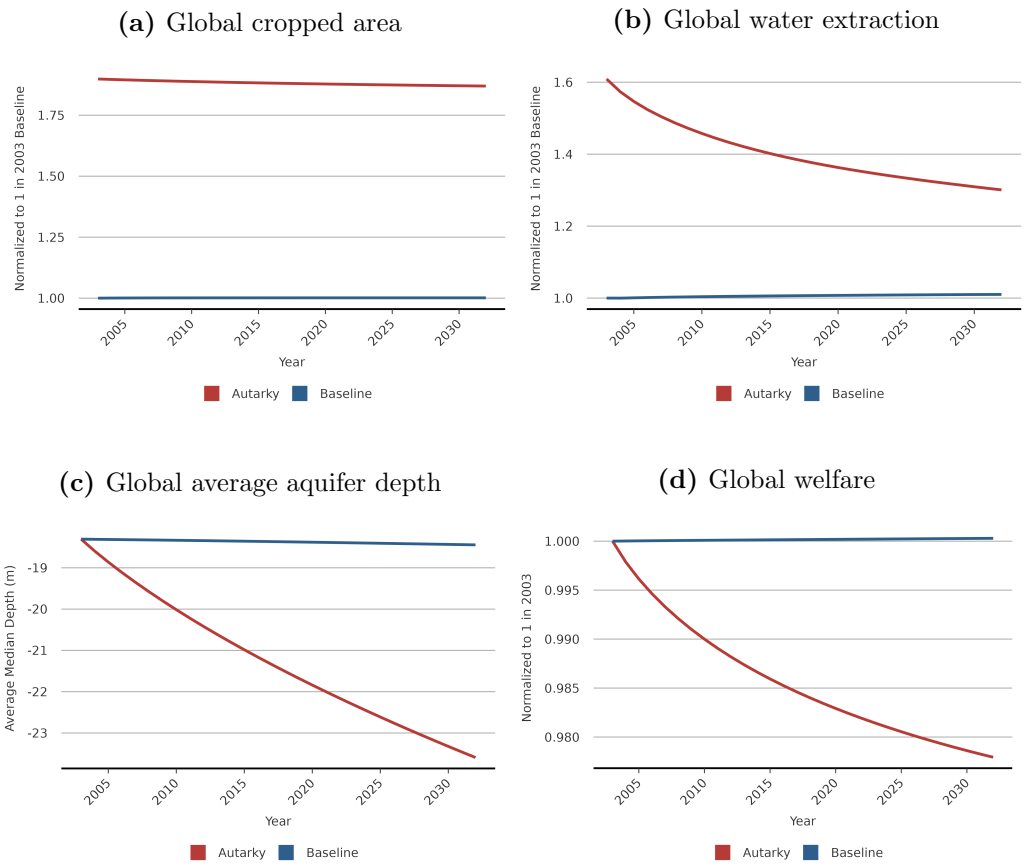

Notes: Each panel shows model simulated output under the baseline calibration (in blue) and a counterfactual autarky simulation (in red) in which trade costs for all crops and all country pairs are infinite.

Figure 8: Relative Water Extraction Under Autarky - By Initial Water Availability

(a) Global Distribution

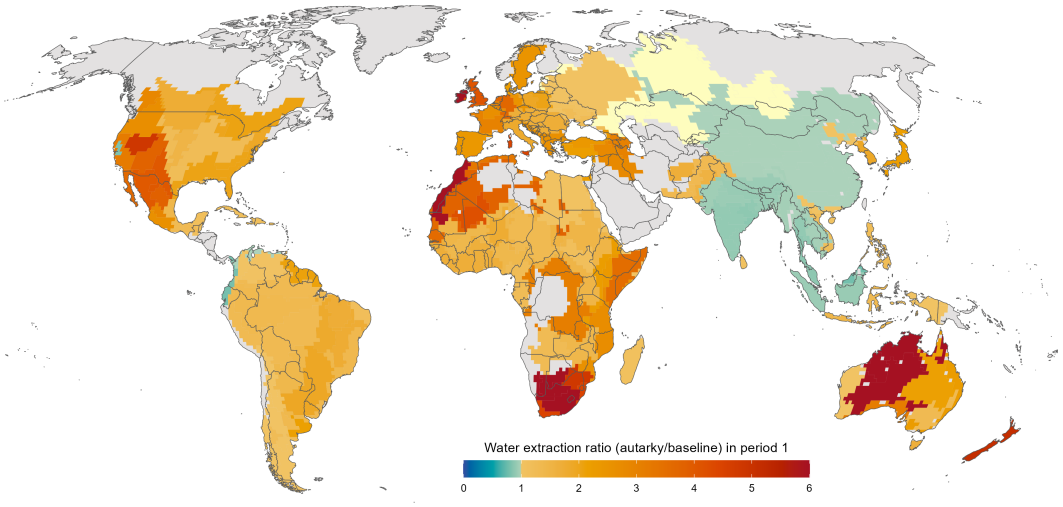

(b) Heterogeneity by Initial Water Resource Characteristics

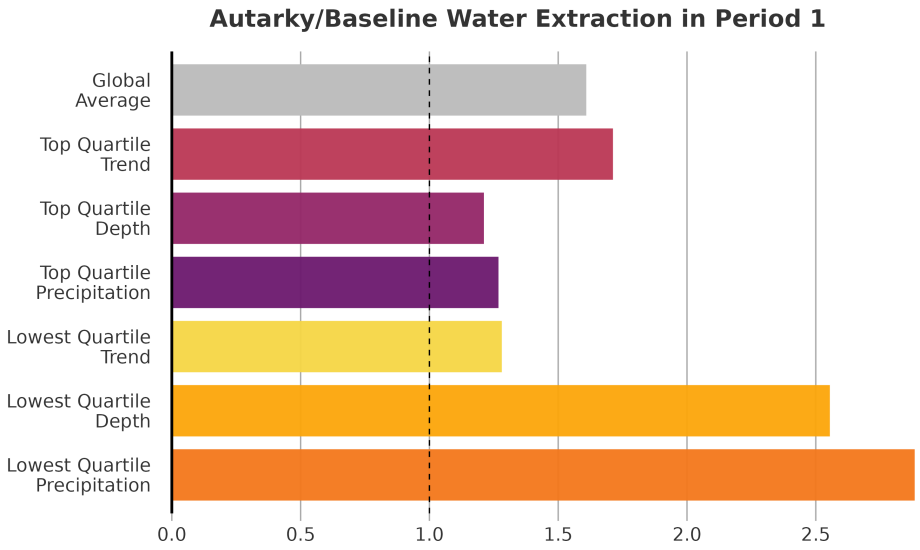

Notes: Map shows the ratio of water extraction under autarky to that under the baseline simulation, at the aquifer level. Ratios above one indicate water extraction rates in autarky that exceed those simulated under baseline agricultural production and trade.

Figure 9: Simulated Trends in Water Tables Over Time

(a) Baseline

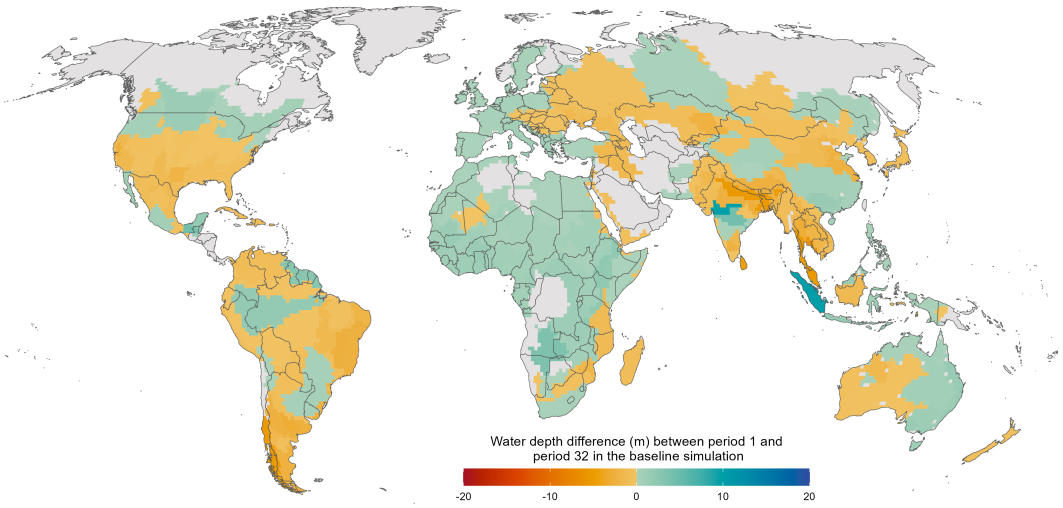

(b) Autarky

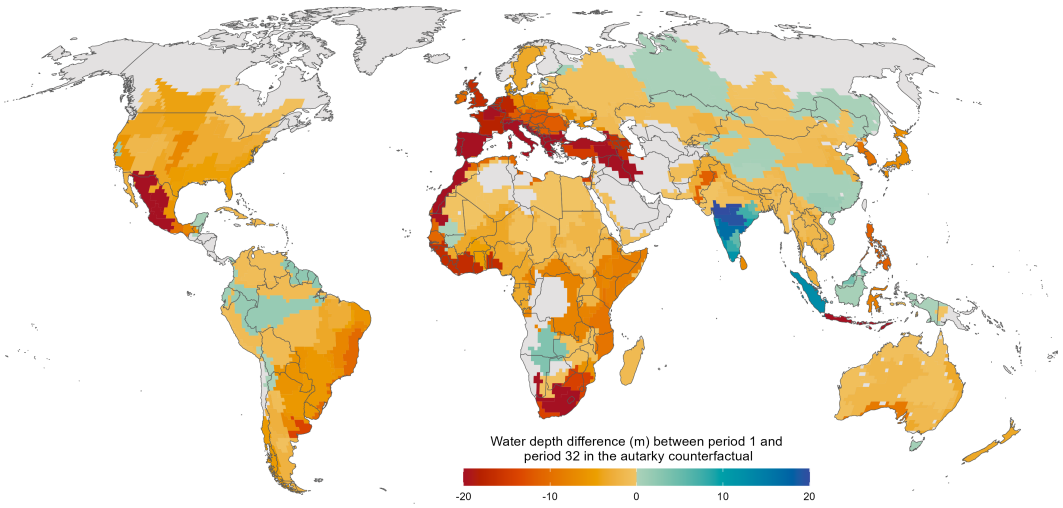

Notes: Map shows the projected 30-year change in water table depth in the baseline simulation in panel (a) and in the autarky counterfactual in panel (b). Changes in water table depth depend on water extraction, recharge, and the specific yield of the soil in each region. Calibrated recharge is held fixed across years and counterfactuals, and Figure 8 shows the ratio of extraction in autarky versus the baseline.

Figure 10: Simulated Impact of Agricultural Trade on Water and Welfare in Select Regions Over Time

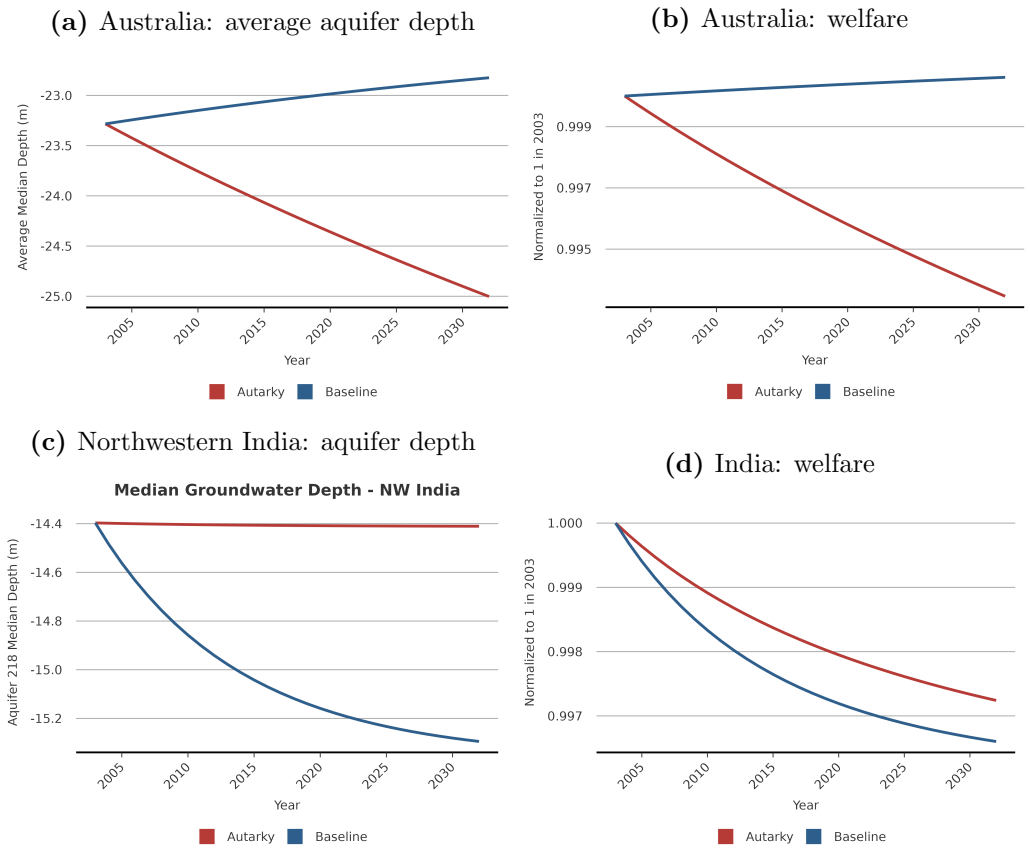

Notes: Each panel shows model simulated output under the baseline calibration (in blue) and a counterfactual autarky simulation (in red) in which trade costs for all crops and all country pairs are infinite. Panels (a) and (b) aggregate all aquifers and fields across Australia, a large importer of water-intensive crops with low groundwater levels. Panel (c) shows the pattern in the most rapidly depleting aquifer in India, a large exporter of water-intensive crops, and panel (d) shows welfare in the aggregate in India, which has a qualitatively similar pattern to the aquifer shown in panel (c) across the country on average.

Figure 11: Nominal Rates of Assistance (NRA) and Simulated Water Extraction Under Alternative Policy Simulations

(a) Changes in NRAs: Uruguay Round

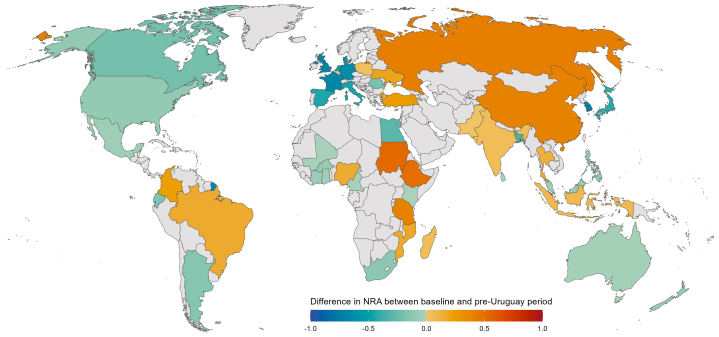

(b) Changes in NRAs: Removing current distortions

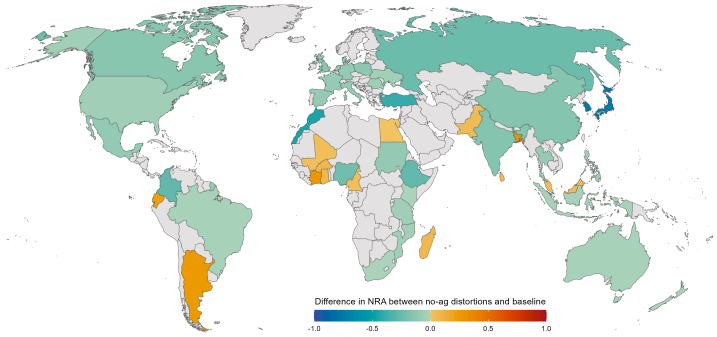

(c) Changes in Water Extraction: Uruguay Round

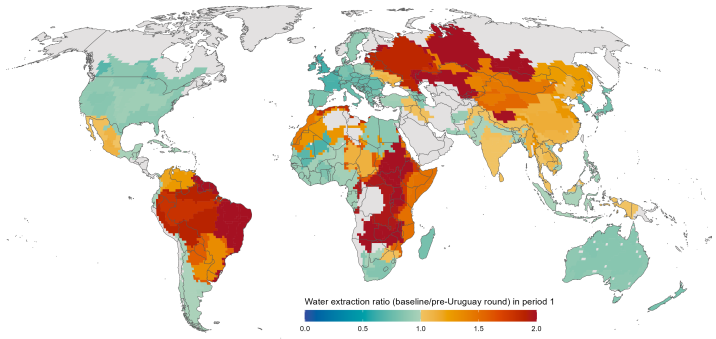

(d) Changes in Water Extraction: Removing current distortions

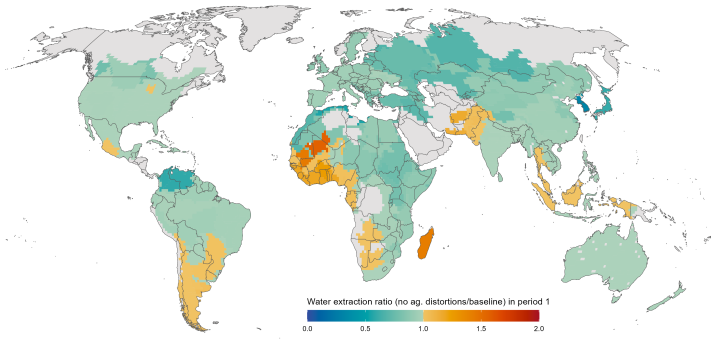

Notes: Panels (a)-(b) show the difference in country-level average Nominal Rates of Assistance (NRA) between (a) after and before the Uruguay Round of World Trade Organization negotiations, which concluded in 1994 and dramatically liberalized agricultural production and trade (negative values indicate that the Uruguay Round lowered subsidies to agriculture); and (b) after versus before removing all baseline (2009) period agricultural distortions (negative values indicate that removing current distortions would lower agricultural subsidies to agriculture). Panels (c)-(d) show global water extraction on impact under these two counterfactual simulations.

Appendix – For Online Publication

A Data Appendix

We compile 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. This compilation includes data on water resources, on agricultural land use, production, and productivity, and on agricultural policy.

A.1 Spatial data on global water resources

Our analysis of the spatial allocation of water resources draws primarily on two scientific datasets that provide information on levels and trends of water stocks 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).[34] 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.[35] 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.[36] 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

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

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

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

area is 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 described in Section 4 of the main text. 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.[37]

Trends in total water availability

The second key water dataset in this paper contains globally gridded information on *changes* in the total volume of local water resources over the period from 2003 through 2016.[38] 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), among many other papers, show that the time series changes in regional mass uncovered by GRACE primarily consist of changes in local water content.

The GRACE data provide month-by-month changes in total water storage (Δ TWS) measured in centimeters of equivalent water height[39] for 41,168 equal-area parcels of the Earth, which measure approximately $1° \times 1°$at the equator.[40] Because the gravity-derived data infers the change in *all* water contained in a given region, they represent the combined evolution of groundwater, surface water, soil moisture, snowpack, and ice over time (though the latter two components are less relevant over the arable lands that are the focus of this paper). These data have been used widely in the scientific literature to study topics ranging from detecting changing water availability (Carleton et al., 2024; Rodell et al., 2018; Richey et al., 2015), assessing ice and glacier melt (Chen et al., 2013), and quantifying changing ocean currents (Wahr et al., 2002), among many other applications. However, its use has been much more limited in economics economics.[41] Throughout our analysis, we

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

[38]This time period corresponds to the original satellite mission from the Gravity Recovery and Climate Experiment (GRACE). A follow-on mission continues to extend this time frame, but we limit our analysis to the original mission for consistency with the years that overlap with other agricultural data used throughout the analysis.

[39]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.

[40]We specifically use the Goddard Space Flight Center's "mascon" solutions, which are equal-area. Other solutions, such as the mascons proceessed by NASA's Jet Propulsion Laboratory, are equal-angle.

[41]The two exceptions to this that we are aware of are recent work by Taylor (2022) and a related

use the raw monthly ΔTWS observations to compute grid-cell specific trends in TWS (in centimeters of equivalent height lost or gained per year) using time series regressions of ΔTWS on day-of-sample, including monthly fixed effects to remove the role of seasonality.

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 in time. 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 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 two datasets relevant to calibrating the law of motion for local water stocks in the model. These include global spatial data on soil type from Hengl et al. (2017) and specific yield by soil type from Loheide et al. (2005). Specific yield is the volume of water that can be drained from porous media by gravity, relative to the total volume of the media. This value, between 0 and 1, indicates the space available in the soil for water mass to be gained or lost. Thus, we use soil type information from Hengl et al. (2017) with specific yield by soil type from Loheide et al. (2005) to estimate the average specific yield at grid cell level across the globe. This then enables us to map changes in the volume of water in a given location into changes in water table depth in model simulations (see Equation 9).

A.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

preceding paper by Carleton (2021).

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.[42] 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 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 by the crop per ton of output. To estimate these values, Mekonnen and Hoekstra (2011) trace 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

[42] 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, to our knowledge, available.

inputs such as nitrogen, and the Monfreda, Ramankutty and Foley (2008) dataset on land use. The quantitative findings from the Mekonnen and Hoekstra (2011) paper, which indicate that over 90% of humanity's water consumption is dedicated to agricultural production as detailed in the main text in Fact 2, are corroborated in a recent review article by d'Odorico et al. (2019). This review states that multiple scientific papers have concluded that approximately 90% of human consumptive use of water arises from agricultural production.

We combine the Mekonnen and Hoekstra (2011) estimates 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 current analysis is that we cannot account for changes in water use arising from different techniques of growing a given crop, or other differences in crop-specific water intensity of production across space.

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 shaping global agricultural trade patterns, and thus 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 percent-

age 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 2, and to calibrate the output market distortions in the model in Section 3.

Table A.1: Global geospatial and administrative datasets used throughout the analysis

	Source	Scale/resolution	Process
Hydrology & geology			
Water table depth	Fan et al. (2013)	1 km grid	Hydrological model used to interpolate between observations from 1.6 million wells
Change in total water storage	GRACE	~1 degree equal-area grid	Satellite-based measurements of anomalies in Earth's gravitational field used to infer changes in water mass
Precipitation & temperature	Sheffield et al. (2006)	0.25 degree grid	Reanalysis product combining a climate model with observational records
Presence of surface water	Pekel et al. (2016)	30 m grid	Machine learning classification of over 3 million Landsat satellite images
Soil type	Hengl et al. (2017)	250 m grid	Machine learning predictions trained on 150,000 soil samples
Specific yield	Loheide et al. (2005)	14 soil types	Literature meta-analysis
Water intensity of production	Mekonnen and Hoekstra (2011)	126 crops	Agronomic and hydrologic modeling using a wide array of input data on climate, soil, and farmer inputs and techniques
Aquifer boundaries	WHYMAP	37 aquifers	Compilation of national, regional and global groundwater sources
Agriculture			
Potential crop yields	FAO GAEZ	~10 km grid × 38 crops	Agronomic modeling using a wide array of climate, soil, and geography input datasets
Agricultural land use	Monfreda et al. (2008)	~10 km grid × 178 crops	Census data combined with remote sensing based downscaling
Area equipped for irrigation	Global Map of Irrigation Areas	~10 km grid	Administrative records combined with geospatial and remote sensing based downscaling
Agricultural production, prices, & trade	FAOSTAT & Comtrade	~200 crops × >200 countries	Direct reporting in administrative records
Policy & other human activity			
Distortions to agricultural incentives	World Bank	80 crops × 82 countries	Direct reporting in administrative records combined with a small open economy model
Nighttime luminosity	U.S. Air Force Defense Meteorological Satellite Program	~1 km grid	Satellite-derived and calibrated imagery of nighttime lights

B Theoretical Appendix

B.1 Derivations

B.1.1 Optimal consumption bundle

We wish to verify equation (10). The representative consumer of country i chooses C_{it}^o and $\{C_{jit}^k\}$ every period to maximize

$$U_{it} = C_{it}^o + \zeta_i \ln C_{it}$$

with the aggregators

$$C_{it} = \left[\sum_{k\in\mathcal{K}} \left(\zeta_i^k\right)^{1/\kappa} \left(C_{it}^k\right)^{\frac{\kappa-1}{\kappa}}\right]^{\frac{\kappa}{\kappa-1}}$$
$$C_{it}^k = \left[\sum_{j\in\mathcal{I}} \left(\zeta_{ji}^k\right)^{1/\sigma} \left(C_{jit}^k\right)^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}}$$

subject to the budget constraint

$$Y_{it} = C_{it}^o + \sum_k \sum_j \delta_{ji}^k \tau_{jt}^k p_{jt}^k C_{jit}^k.$$

We solve the problem nest-by-nest:

(i) what is the optimal cross-country bundle $\{C_{jit}^k\}_j$ for a given crop-specific index C_{it}^k?

(ii) what is the optimal cross-crop bundle $\{C_{it}^k\}_k$ for a given agricultural index C_{it}?

(iii) what is the optimal expenditure split between the outside good C_{it}^o and the agricultural index C_{it}?

The usual CES derivations give us the following relationships for (i) and (ii):

$$C_{jit}^k = \zeta_{ji}^k \left(\delta_{ji}^k \tau_{jt}^k p_{jt}^k\right)^{-\sigma} \left(P_{it}^k\right)^{\sigma} C_{it}^k \quad \text{with price index } P_{it}^k = \left[\sum_n \zeta_{ni}^k \left(\delta_{ni}^k \tau_{nt}^k p_{nt}^k\right)^{1-\sigma}\right]^{\frac{1}{1-\sigma}}$$
$$C_{it}^k = \zeta_i^k \left(P_{it}^k\right)^{-\kappa} (P_{it})^{\kappa} C_{it} \quad \text{with price index } P_{it} = \left[\sum_k \zeta_i^k \left(P_{it}^k\right)^{1-\kappa}\right]^{\frac{1}{1-\kappa}}.$$

The outermost nest (iii) is simply the solution to

$$\max_{C_{it}^o, C_{it}} \; C_{it}^o + \zeta_i \ln C_{it} \quad \text{s.t. } Y_{it} = C_{it}^o + P_{it} C_{it},$$

which is $P_i C_i = \zeta_i$ and $C^o_{it} = Y_{it} - \zeta_i$. Putting everything together:

$$\begin{aligned}
C^k_{jit} &= \zeta^k_{ji} \left(\delta^k_{ji}\tau^k_{jt}p^k_{jt}\right)^{-\sigma} \left(P^k_{it}\right)^{\sigma} C^k_{it} \\
&= \zeta^k_{ji} \left(\delta^k_{ji}\tau^k_{jt}p^k_{jt}\right)^{-\sigma} \left(P^k_{it}\right)^{\sigma} \zeta^k_i \left(P^k_{it}\right)^{-\kappa} (P_{it})^{\kappa} C_{it} \\
&= \zeta^k_{ji} \left(\delta^k_{ji}\tau^k_{jt}p^k_{jt}\right)^{-\sigma} \left(P^k_{it}\right)^{\sigma} \zeta^k_i \left(P^k_{it}\right)^{-\kappa} (P_{it})^{\kappa-1} \zeta_i \\
&= \zeta^k_{ji} \left(\delta^k_{ji}\tau^k_{jt}p^k_{jt}\right)^{-\sigma} \left(P^k_{it}\right)^{\sigma-1} \zeta^k_i \left(P^k_{it}\right)^{1-\kappa} (P_{it})^{\kappa-1} \zeta_i \\
&= \frac{\zeta^k_{ji} \left(\delta^k_{ji}\tau^k_{jt}p^k_{jt}\right)^{-\sigma}}{\left(P^k_{it}\right)^{1-\sigma}} \frac{\zeta^k_i \left(P^k_{it}\right)^{1-\kappa}}{(P_{it})^{1-\kappa}} \zeta_i,
\end{aligned}$$

which, after reversing the order of terms and substituting in the definitions of the price indices, gives us (10).

B.1.2 Optimal labor allocation and water use

For a given crop k, the farmer solves the labor allocation problem

$$\max_{H \in [0,1]} Q^{fk}(\omega)[H] = A^{fk}(\omega) H^{\alpha} \left[N^{fk}(H)\right]^{1-\alpha}$$

where $N^{fk}(H) \equiv \min\left\{1, \frac{A^w(D_{q(f)})[1-H]}{\phi^k}\right\}$ denotes the "natural capital" (i.e., the effective bundle of land and groundwater) the farmer can use to grow crop k if he allocates fraction H of his labor to tending the crop (hence $1-H$ to water extraction). An equivalent formulation of $N^{fk}(H)$ is

$$N^{fk}(H) = \begin{cases} 1 & \text{if } H \leq 1 - \frac{\phi^k}{A^w(D_{q(f)})}, \\ \frac{A^w(D_{q(f)})}{\phi^k}(1-H) & \text{if } H \geq 1 - \frac{\phi^k}{A^w(D_{q(f)})}. \end{cases}$$

Note that only in the first case would the farmer be irrigating his entire parcel.

We wish to derive the farmer's optimal policy function for labor, $H^{fk}(\omega)$, which in turn will determine how much water he extracts. We can prove almost immediately that the labor allocated to extracting groundwater will be no more than $\frac{\phi^k}{A^w(D_{q(f)})}$ because land and water are used in fixed proportions.

Lemma B.1. $H^{fk}(\omega) \geq 1 - \frac{\phi^k}{A^w(D_{q(f)})}$.

Proof. $N^{fk}(H)$ is constant but $A^{fk}(\omega)H$ is strictly increasing—hence $Q^{fk}(\omega)[H]$ is strictly increasing—for any $H < 1 - \frac{\phi^k}{A^w(D_{q(f)})}$, so the farmer's optimal H cannot be in that range. □

Now consider the range over which $H > 1 - \frac{\phi^k}{A^w(D_{q(f)})}$ and take the derivative of the farmer's objective with respect to H:

$$\begin{aligned}
&\frac{d}{dH}\left\{A^{fk}(\omega)H^{\alpha}\left[N^{fk}(H)\right]^{1-\alpha}\right\} \\
&\qquad = A^{fk}(\omega)\left[\alpha\left(\frac{N^{fk}(H)}{H}\right)^{1-\alpha} + (1-\alpha)(N^{fk})'(H)\left(\frac{H}{N^{fk}(H)}\right)^{\alpha}\right] \\
&\qquad = A^{fk}(\omega)\left(\frac{A^w(D_{q(f)})}{\phi^k}\right)^{1-\alpha}\left[\alpha\left(\frac{1-H}{H}\right)^{1-\alpha} - (1-\alpha)\left(\frac{H}{1-H}\right)^{\alpha}\right].
\end{aligned}$$

Setting this derivative equal to zero yields $H^{fk}(\omega) = \alpha$ over that range. Combining this result with Lemma B.1, it follows that the optimal policy function for labor is

$$H^{fk} = \begin{cases} 1 - \frac{\phi^k}{A^w(D_{q(f)})} & \text{if } \frac{\phi^k}{A^w(D_{q(f)})} \leq 1-\alpha \\ \alpha & \text{if } \frac{\phi^k}{A^w(D_{q(f)})} > 1-\alpha, \end{cases} \tag{B.1}$$

where we've dropped the argument ω identifying the parcel because it proved irrelevant.

Accordingly, we can summarize the output of crop k by a farmer on parcel ω of field f as $A^{fk}(\omega)M(\phi^k, D_{q(f)})$, where

$$M(\phi^k, D_{q(f)}) = \begin{cases} \left(1 - \frac{\phi^k}{A^w(D_{q(f)})}\right)^{\alpha} & \text{if } \frac{\phi^k}{A^w(D_{q(f)})} \leq 1-\alpha \\ \tilde{\alpha}\left(\frac{A^w(D_{q(f)})}{\phi^k}\right)^{1-\alpha} & \text{if } \frac{\phi^k}{A^w(D_{q(f)})} > 1-\alpha \end{cases} \tag{B.2}$$

with $\tilde{\alpha} = \alpha^{\alpha}(1-\alpha)^{1-\alpha}$. So output is always decreasing in the water requirement ϕ^k and the water table depth D. Note that this formulation allows for output to approach zero (as depth approaches infinity) without any discontinuities or edge cases.

Likewise, we see that extraction by a farmer on field f cultivating crop k is

$$x^{fk} = \begin{cases} \phi^k & \text{if } \frac{\phi^k}{A^w(D_{q(f)})} \leq 1-\alpha \\ (1-\alpha)A^w(D_{q(f)}) & \text{if } \frac{\phi^k}{A^w(D_{q(f)})} > 1-\alpha. \end{cases} \tag{B.3}$$

B.1.3 Crop-specific cropped area shares (by field)

For convenience, assume i means $i(f)$ and q means $q(f)$ unless otherwise specified. We wish to verify equation (11). The probability that crop k is grown on any given parcel ω of field f (where $k = o$ means the farmer works in the outside sector and leaves his parcel fallow) is the probability that the revenue from k is highest. Note

that the c.d.f. of $r_t^{fk}(\omega)$ is

$$F_t^{fk}(r) = \exp\left\{-\gamma\left(\frac{r}{\tau_{it}^k p_{it}^k A^{fk} M(\phi^k, D_{qt})}\right)^{-\theta}\right\}$$

which follows directly from $A^{fk}(\omega)$ being distributed Fréchet and the definition $r_t^{fk}(\omega) \equiv \tau_{it}^k p_{it}^k A^{fk}(\omega) M(\phi^k, D_{qt})$. Note, too, that the same can be said for the outside good $k = o$ if we set $p_{it}^o = 1$, $A^{fo} = A_i^o$, and $\phi^o = 0$. Thus, we compute

$$\begin{aligned}
&\mathbb{P}\left\{r_t^{fk}(\omega) = \max\{A_i^o(\omega), r_t^{f1}(\omega), \ldots, r_t^{fK}(\omega)\}\right\} \\
&\quad= \int_0^\infty \prod_{\ell \neq k} F_t^{f\ell}(r)\, d\left[1 - F_t^{fk}(r)\right] \\
&\quad= \int_0^\infty \prod_{\ell \neq k} F_t^{f\ell}(r) \times \left[-f_t^{fk}(r)\right] dr \\
&\quad= \int_0^\infty \prod_{\ell \in \mathcal{K} \cup \{o\}} \exp\left\{-\gamma\left(\frac{r}{\tau_{it}^\ell p_{it}^\ell A_i^{f\ell} M(\phi^k, D_{qt})}\right)^{-\theta}\right\} \frac{-\gamma\theta}{\left(\tau_{it}^k p_{it}^k A^{fk} M(\phi^k, D_{qt})\right)^{-\theta}} r^{-\theta-1}\, dr \\
&\quad= \left(\tau_{it}^k p_{it}^k A^{fk} M(\phi^k, D_{qt})\right)^\theta \int_0^\infty -\exp\left\{-\gamma r^{-\theta} \sum_{\ell \in \mathcal{K} \cup \{o\}} \left(\tau_{it}^\ell p_{it}^\ell A_i^{f\ell} M(\phi^k, D_{qt})\right)^\theta\right\} \gamma\theta r^{-\theta-1}\, dr \\
&\quad= \left(\tau_{it}^k p_{it}^k A^{fk} M(\phi^k, D_{qt})\right)^\theta \int_0^\infty -\exp\left\{-\gamma\left[(A_i^o)^\theta + \Pi_t^f\right] r^{-\theta}\right\} \gamma\theta r^{-\theta-1}\, dr \\
&\quad= \frac{\left(\tau_{it}^k p_{it}^k A^{fk} M(\phi^k, D_{qt})\right)^\theta}{(A_i^o)^\theta + \Pi_t^f} \int_0^\infty d\left[1 - F_t^f(r)\right] \\
&\quad= \frac{\left(\tau_{it}^k p_{it}^k A^{fk} M(\phi^k, D_{qt})\right)^\theta}{(A_i^o)^\theta + \Pi_t^f}
\end{aligned}$$

as claimed, where

$$1 - F_t^f(r) \equiv 1 - \exp\left\{-\gamma\left[(A_i^o)^\theta + \Pi_t^f\right] r^{-\theta}\right\}.$$

is the probability that the revenue for at least one use of a given parcel on field f (including leaving it fallow) exceeds r.

B.1.4 Expected productivity given crop selection

Continue to assume that i means $i(f)$ and q means $q(f)$ unless otherwise specified. We wish to show that

$$\mathbb{E}\left[A^{fk}(\omega)\middle| r_t^{fk}(\omega) = \max\{A_i^o(\omega), r_t^{f1}(\omega), \ldots, r_t^{fK}(\omega)\}\right] = A^{fk}\left(\pi_t^{fk}\right)^{-1/\theta}.$$

To compute this conditional expected value, we need to derive the corresponding conditional probability distribution. The probability that crop k is grown on field f with productivity a or less given that it is the profit-maximizing crop choice for field f can be decomposed as (i) the joint probability that crop k earns return at or below $r_t^k(a) \equiv \tau_{it}^k p_{it}^k a M(\phi^k, D_{qt})$ and is the profit-maximizing crop choice for f divided by (ii) the unconditional probability of being the profit-maximizing crop, π_t^{fk}. Then,

$$
\begin{aligned}
&\mathbb{P}\left\{ r_t^{fk}(\omega) < r_t^k(a) \middle| r_t^{fk}(\omega) = \max\{A_i^o(\omega), r_t^{f1}(\omega), \ldots, r_t^{fk}(\omega)\} \right\} \\
&\quad = \frac{1}{\pi_t^{fk}} \int_0^{r_t^k(a)} \prod_{\ell \neq k} F_t^{f\ell}(r)\, d\left[1 - F_t^{fk}(r)\right] \\
&\quad = \frac{1}{\pi_t^{fk}} \int_0^{r_t^k(a)} \prod_{\ell \neq k} F_t^{f\ell}(r) \times \left[-f_t^{fk}(r)\right]\, dr \\
&\quad = \frac{1}{\pi_t^{fk}} \int_0^{r_t^k(a)} \prod_{\ell \in \mathcal{K} \cup \{o\}} \exp\left\{ -\gamma \left(\frac{r}{\tau_{it}^\ell p_{it}^\ell A_i^{f\ell} M(\phi^k, D_{qt})} \right)^{-\theta} \right\} \frac{-\gamma\theta}{\left(\tau_{it}^k p_{it}^k A^{fk} M(\phi^k, D_{qt})\right)^{-\theta}} r^{-\theta-1}\, dr \\
&\quad = \frac{\left(\tau_{it}^k p_{it}^k A^{fk} M(\phi^k, D_{qt})\right)^\theta}{\pi_t^{fk}} \int_0^{r_t^k(a)} -\exp\left\{ -\gamma r^{-\theta} \sum_{\ell \in \mathcal{K} \cup \{o\}} \left(\tau_{it}^\ell p_{it}^\ell A_i^{f\ell} M(\phi^k, D_{qt})\right)^\theta \right\} \gamma\theta r^{-\theta-1}\, dr \\
&\quad = \frac{\left(\tau_{it}^k p_{it}^k A^{fk} M(\phi^k, D_{qt})\right)^\theta}{\pi_t^{fk}} \int_0^{r_t^k(a)} -\exp\left\{ -\gamma \left[(A_i^o)^\theta + \Pi_t^f \right] r^{-\theta} \right\} \gamma\theta r^{-\theta-1}\, dr \\
&\quad = \frac{1}{\pi_t^{fk}} \left[\frac{\left(\tau_{it}^k p_{it}^k A^{fk} M(\phi^k, D_{qt})\right)^\theta}{(A_i^o)^\theta + \Pi_t^f} \right] \left[1 - F_t^f\left(r_t^k(a)\right) \right] \\
&\quad = 1 - F_t^f\left(r_t^k(a)\right).
\end{aligned}
$$

The expected value of $r_t^{fk}(\omega)$ against this conditional distribution is

$$
\begin{aligned}
&\mathbb{E}\left[r_t^{fk}(\omega) \middle| r_t^{fk}(\omega) = \max\{A_i^o(\omega), r_t^{f1}(\omega), \ldots, r_t^{fK}(\omega)\} \right] \\
&\qquad = \int_0^\infty r\, d\left[1 - F_t^f(r)\right] \\
&\qquad = \int_0^\infty -r\gamma\theta \left[(A_i^o)^\theta + \Pi_t^f \right] r^{-\theta-1} \exp\left\{ -\gamma \left[(A_i^o)^\theta + \Pi_t^f \right] r^{-\theta} \right\} dr,
\end{aligned}
$$

and now with the variable transformation $x \equiv \gamma \left[(A_i^o)^\theta + \Pi_t^f\right] r^{-\theta}$,

$$
\begin{aligned}
&= \int_0^\infty r \left(\frac{dx}{dr}\right) \exp\{-x\}\, dr \\
&= \int_0^\infty \left(\frac{x}{\gamma\left[(A_i^o)^\theta + \Pi_t^f\right]}\right)^{-1/\theta} \exp\{-x\}\, dx \\
&= \gamma^{1/\theta}\left[(A_i^o)^\theta + \Pi_t^f\right]^{1/\theta} \int_0^\infty x^{-1/\theta} \exp\{-x\}\, dx \\
&= \gamma^{1/\theta}\left[(A_i^o)^\theta + \Pi_t^f\right]^{1/\theta} \Gamma\left(\frac{\theta-1}{\theta}\right) \\
&= \left[(A_i^o)^\theta + \Pi_t^f\right]^{1/\theta}.
\end{aligned}
$$

Now transforming from revenue to productivity:

$$
\begin{aligned}
\mathbb{E}\left[A_t^{fk}(\omega) \middle| r_t^{fk}(\omega) = \max\{A_i^o(\omega), r_t^{f1}(\omega), \ldots, r_t^{fk}(\omega)\}\right] &= \frac{\left[(A_i^o)^\theta + \Pi_t^f\right]^{1/\theta}}{\tau_{it}^k p_{it}^k M(\phi^k, D_{qt})} \\
&= A^{fk} \frac{\left[(A_i^o)^\theta + \Pi_t^f\right]^{1/\theta}}{\tau_{it}^k p_{it}^k A^{fk} M(\phi^k, D_{qt})} \\
&= A^{fk}\left(\pi_t^{fk}\right)^{-1/\theta}
\end{aligned}
$$

which completes the proof.

B.1.5 Robustness: Nested CES agricultural technology

Following Boppart et al. (2019), who reject Cobb-Douglas technology in the agricultural sector in favor of nested CES technologies, we consider the technology (omitting sub- and superscripts for convenience)

$$
Q = A\left[\alpha^{\frac{1}{\xi_H}} H^{\frac{\xi_H - 1}{\xi_H}} + (1-\alpha)^{\frac{1}{\xi_H}} \left(L^{\frac{\xi_M - 1}{\xi_M}} + \left(\frac{G}{\phi^k}\right)^{\frac{\xi_M - 1}{\xi_M}}\right)^{\frac{\xi_M}{\xi_M - 1}\frac{\xi_H - 1}{\xi_H}}\right]^{\frac{\xi_H}{\xi_H - 1}}. \quad \text{(B.4)}
$$

One can verify that in the limit as $\xi_H \to 1$ and $\xi_M \to 0$, this becomes the technology in (4).

A crucial step in characterizing the equilibrium of the model is showing that agricultural output on a given parcel can be written as the product of its idiosyncratic TFP draw and a function M that is independent of the realization of that

TFP draw (see Section B.1.2). From the farmer's first-order conditions with technology (B.4), one can verify that

$$Q_t^{fk}(\omega) = A^{fk}(\omega)\tilde{M}(\phi^k, D_{q(f)t})$$

with

$$\tilde{M}(\phi^k, D_{q(f)t}) = \left[\alpha^{\frac{1}{\xi_H}}(\tilde{H})^{\frac{\xi_H-1}{\xi_H}} + (1-\alpha)^{\frac{1}{\xi_H}}\left(1+\left(\frac{A^w(D_{q(f)t})[1-\tilde{H}]}{\phi^k}\right)^{\frac{\xi_M-1}{\xi_M}}\right)^{\frac{\xi_M}{\xi_M-1}\frac{\xi_H-1}{\xi_H}}\right]^{\frac{\xi_H}{\xi_H-1}},$$

where $\tilde{H}$ is implicitly defined by the condition

$$\tilde{H} = \frac{\alpha}{1-\alpha}\left[\left(\frac{A^w(D_{q(f)t})}{\phi^k}\right)^{\frac{\xi_M-1}{\xi_M}}(1-\tilde{H})^{-\frac{1}{\xi_M}}\left(1+\left[\frac{A^w(D_{q(f)t})[1-\tilde{H}]}{\phi^k}\right]^{\frac{\xi_M-1}{\xi_M}}\right)^{\frac{1}{\xi_M-1}}\right].$$

So one can endow farmers with the technology in (B.4) and define an equilibrium as before, just in terms of $\tilde{M}$. But doing so would require us to find the root of the preceding nonlinear equation every time the depth of an aquifer changed, which will happen for every aquifer in every period along the equilibrium path. Moreover, to the best of our knowledge, there is no globally comprehensive data with which to estimate the elasticities ξ_H and ξ_M, nor a natural benchmark to which they can be calibrated. As such, we chose the Leontief specification presented in the main text.

B.2 Existence and uniqueness

The goal of this section is to characterize under what conditions the competitive equilibrium given by Definition 1 exists and is unique. To that end, we'll proceed by defining and characterizing two sub-equilibria: a *trade equilibrium*, which takes the vector of aquifer depths as given, and a *steady-state equilibrium*, which imposes constant depths. Once we've characterized those two sub-equilibria, we'll characterize the *full dynamic competitive equilibrium* given by Definition 1. Because agricultural policy distortions are exogenous and act proportionally on prices, they would impose no additional complications on the arguments below, so they've been ignored to save on notation.

B.2.1 ... of the trade equilibrium

A *trade equilibrium* asks only how goods and factors will be allocated today given depths $\mathbf{D}$; it does not consider how the allocation today will affect depths tomorrow. A formal definition follows.

Definition B.1. Given an arbitrary vector of groundwater depths, $\mathbf{D}$, a *trade equilibrium* is a vector of consumption, $\{C_{ji}^k\}$, output, $\{Q_i^k\}$, prices, $\{p_i^k\}$, and shares, $\{\pi^{fk}\}$, such that equations (10), (11), (12), and (13) hold.

With this reduction in scope, we've recovered a static neoclassical trade model. The presence of the outside good means that this trade block falls outside the class of gravity models that have been characterized by Allen et al. (2020).[43] Instead, we'll follow Alvarez and Lucas (2007) in using the *aggregate excess demand function* $\mathbf{z}(\tilde{\mathbf{p}})$:

$$z_{ik}(\tilde{\mathbf{p}}) = \sum_{j\in\mathcal{I}} \delta_{ij}^k (p^o \zeta_j) \frac{\zeta_j^k \left(P_j^k\right)^{1-\kappa}}{\sum_{\ell\in\mathcal{K}} \zeta_j^\ell \left(P_j^\ell\right)^{1-\kappa}} \frac{\zeta_{ij}^k \left(\delta_{ij}^k p_i^k\right)^{-\sigma}}{\sum_{n\in\mathcal{I}} \zeta_{nj}^k \left(\delta_{nj}^k p_n^k\right)^{1-\sigma}}$$
$$- \sum_{f\in\mathcal{F}_i} h^f A_i^{fk} M(\phi^k, D_{q(f)}) \left(\frac{\left(p_i^k A_i^{fk} M(\phi^k, D_{q(f)})\right)^\theta}{(p^o A_i^o)^\theta + \sum_{\ell\in\mathcal{K}} \left(p_i^\ell A_i^{f\ell} M(\phi^\ell, D_{q(f)})\right)^\theta} \right)^{\frac{\theta-1}{\theta}}$$

and

$$z_o(\tilde{\mathbf{p}}) = \sum_i \left(\frac{Y_i}{p^o} - \zeta_i\right) - \sum_i \sum_{f\in\mathcal{F}_i} h^f A_i^o \left(\frac{(p^o A_i^o)^\theta}{(p^o A_i^o)^\theta + \sum_{\ell\in\mathcal{K}} \left(p_i^\ell A_i^{f\ell} M(\phi^\ell, D_{q(f)})\right)^\theta} \right)^{\frac{\theta-1}{\theta}}$$

where $\tilde{\mathbf{p}} \equiv [\mathbf{p}, p^o]$, with p^o being the price of the outside good (which we had been normalizing to one), and Y_i is total income $Y_i = p^o Q_i^o + \sum_k p_i^k Q_i^k$. By Mas-Colell et al. (1995, Ch.17), the equilibrium exists if

1. $\mathbf{z}$ is continuous;
2. $\mathbf{z}$ is homogeneous of degree zero in $\tilde{\mathbf{p}}$;
3. $\tilde{\mathbf{p}} \cdot \mathbf{z} = 0$ for all strictly positive price vectors (Walras' law);
4. there is a $\underline{z} > 0$ such that $z_\ell(\tilde{\mathbf{p}}) > -\underline{z}$ for every commodity ℓ and all p;
5. if $\tilde{\mathbf{p}}^n \to \tilde{\mathbf{p}}$, where $\tilde{\mathbf{p}} \neq 0$ but $\tilde{p}_\ell = 0$ for some ℓ, then

$$\max\{z_1(\tilde{\mathbf{p}}^n), \ldots, z_{IK+1}(\tilde{\mathbf{p}}^n)\} \to \infty;$$

[43] Considering the agricultural and outside sectors together, aggregate demand here is not CES, which violates condition C.2 of Allen et al. (2020). One may consider only the agricultural sector, which is well-described by an Armington model of trade, but then expenditure on the outside good is isomorphic to an endogenous trade imbalance, which violates condition C.5. Note, too, that we cannot map production in each country a representative good, which violates condition C.3.

and is unique if, in addition,

6. $\frac{\partial z_\ell(\tilde{\mathbf{p}})}{\partial p_{\ell'}} > 0$ for all ℓ, ℓ' with $\ell \neq \ell'$ and all $\tilde{\mathbf{p}} \in \mathbb{R}^{IK+1}_{++}$ (*gross substitutes*).

Here we're leveraging that we can identify the trade equilibrium by just the vector of prices (up to scale): Once we have the prices, the vectors of consumption, output, and land use shares are determined uniquely.

Proposition B.2. *A trade equilibrium exists for any vector of depths* $\mathbf{D}$.

Proof. We just need to check conditions 1-5. We'll do so in order.

1. Continuity is self-evident.

2. Homogeneity of degree zero is also self-evident (multiply every price by the same factor $\lambda > 0$ and confirm that all the multipliers cancel).

3. Let's look at the demand components first. In z_{ik} the dot product gives us $\sum_j p_i^k \delta_{ij}^k C_{ij}^k$. In z_o we get $\sum_i (Y_i - p^o \zeta_i)$. Now the supply terms: in z_{ik} we get $p_i^k Q_i^k$ and in z_o we get $p^o Q_i^o$. So

$$\begin{aligned}
\tilde{\mathbf{p}} \cdot \mathbf{z} &= \sum_{i,k,j} p_i^k \delta_{ij}^k C_{ij}^k + \sum_i (Y_i - p^o \zeta_i) - \sum_{i,k} p_i^k Q_i^k - \sum_i p^o Q_i^o \\
&= \sum_{i,k,j} p_i^k \delta_{ij}^k C_{ij}^k - \sum_i p^o \zeta_i \\
&= \sum_{i,k,j} p_i^k \delta_{ij}^k C_{ij}^k - \sum_i P_i C_i \\
&= 0.
\end{aligned}$$

4. The demand terms are always non-negative, so set them to zero. For each element of the supply terms, take the maximum value permissible (so $\max_f h^f$ and $\max_{fk} A^{fk}$ and so on). Note that the max permissible share is always 1. Since each maximum element is finite, one can construct a finite $\bar{z}$ using these maximums.

5. Note that, defining $\Lambda^{fk} \equiv h^f A^{fk} M(\phi^k, D_{q(f)})$ that is independent of prices for fixed D,

$$\begin{aligned}
\max_{ik} z_{ik}(\tilde{\mathbf{p}}) &\geq \max_{ik} \sum_j \delta_{ij}^k C_{ij}^k - \max_{ik} \sum_{f \in \mathcal{F}_i} \Lambda^{fk} (\pi^{fk})^{(\theta-1)/\theta} \\
&\geq \max_{ik,j} \delta_{ij}^k C_{ij}^k - \max_{ik} \sum_{f \in \mathcal{F}_i} \Lambda^{fk} (\pi^{fk})^{(\theta-1)/\theta} \\
&\geq \max_{ik,j} \delta_{ij}^k C_{ij}^k - \max_{ik} \sum_{f \in \mathcal{F}_i} \Lambda^{fk}
\end{aligned}$$

where we get each inequality because, in order,

(i) we can show $\max_k(A_k - B_k) \geq \max_k A_k - \max_k B_k$ as follows:
   - suppose $k^\delta = \arg\max_k(A_k - B_k)$ and $k^A = \arg\max_k A_k$
   - let $\delta \equiv A_{k^A} - A_{k^\delta}$
   - consider the difference $A_{k^\delta} - B_{k^\delta} = A_{k^\delta} + \delta - (B_{k^\delta} + \delta) = A_{k^A} - (B_{k^\delta} + \delta)$
   - now $B_{k^\delta} + \delta \leq B_{k^A} \leq \max_k B_k$, else k^A would be the arg max of $A_k - B_k$, not k^δ;

(ii) a sum of positive terms is greater than its maximum term;

(iii) $\pi^{fk} \leq 1$ for all fk.

So the task reduces to showing that $\max_{ik,j} C_{ij}^k(\tilde{\mathbf{p}}) \to \infty$ for the given price sequence, where we overload notation (as we're wont to do) by considering $k \in \{\mathcal{K}, o\}$. Consider the definition of C_{ij}^k (for $k \in \mathcal{K}$) term-by-term:

- the leftmost is $p^o\zeta_i$, which is obviously finite whenever p^o is finite
- the middle is bounded between zero and one
- the rightmost is also bounded between zero and one when pre-multiplied by $\delta_{ij}^k p_i^k$, so the rightmost on its own tends to infinity as $p_i^k \to 0$.

And if $p^o \to 0$, it's obvious from the definition of C_i^o that it will tend to infinity.

□

Proposition B.3. *The trade equilibrium is unique for any vector of depths* $\mathbf{D}$.

Proof. We just need to check the gross substitutes condition.[44] This is easy to verify for the outside good with respect to any agricultural commodity price: in the expression for $z_o(\tilde{\mathbf{p}})$, agricultural prices only show up in the denominator of the term being subtracted. Now consider the excess demand for an agricultural commodity, $z_{ik}(\tilde{\mathbf{p}})$. The gross substitutes condition is satisfied with respect to the outside good: if p^o increased, demand for any agricultural commodity would increase (the first term), but supply of every agricultural commodity would decrease (the second term), so excess demand for ik would surely increase. With respect to any agricultural commodity from another country $i' \neq i$, an increase in its price implies a direct increase in demand for ik but no change in its supply, so $z_{ik}(\tilde{\mathbf{p}})$ increases. Finally, consider the direct effect of an increase in price for a commodity from the same country, ik' with $k' \neq k$. By the same logic as before, demand for ik increases as consumers substitute away from the higher price (first term) and supply of ik decreases as domestic producers substitute toward the higher priced ik' (second term). Thus, excess demand for ik increases. □

[44] The expressions for the relevant partial derivatives are easy to derive but repetitive to show, so we omit them here in favor of a verbal argument.

B.2.2 ... of the steady-state equilibrium

A *steady state* is a path along which the depth of each aquifer is constant over time. Since depths were the only variables evolving with any fundamental persistence, one can define a steady state as a trade equilibrium in which the inflows to each aquifer offset the outflows.

Definition B.2. A *steady-state equilibrium* is a vector of consumption, $\{C_{ji}^k\}$, output, $\{Q_i^k\}$, prices, $\{p_i^k\}$, shares, $\{\pi^{fk}\}$, groundwater depths, $\{D_q\}$, and groundwater extractions, $\{X_q\}$, such that

$$(1-\psi)X_q = R_q, \quad \forall q \in \mathcal{Q}$$

and equations (10), (11), (12), (13), and (14) hold.

Proposition B.4. *A unique steady-state equilibrium exists.*

Proof. Propositions B.2–B.3 established that for any vector of depths there exists a unique vector of crop prices, which uniquely determines the rest of the trade allocation. Accordingly, we can define the function

$$X_q(D_q) = \sum_{f \in \mathcal{F}_q} \sum_{k \in \mathcal{K}} x^{fk} h^f \frac{\left(p_{i(f)}^k A^{fk} M(\phi^k, D_q)\right)^\theta}{\left(A_{i(f)}^o\right)^\theta + \sum_{\ell \in \mathcal{K}} \left(p_{i(f)}^\ell A^{f\ell} M(\phi^\ell, D_q)\right)^\theta},$$

where the prices $p_{i(f)}^k$ are those from the corresponding trade equilibrium, so that a steady-state equilibrium is $\bar{\mathbf{D}} = \{\bar{D}_q\}$ such that

$$(1-\psi)X_q(\bar{D}_q) = R_q, \quad \forall q \in \mathcal{Q}.$$

Equivalently, we can define the operator $T : \mathbb{R}^{|\mathcal{Q}|} \mapsto \mathbb{R}^{|\mathcal{Q}|}$ by

$$[T(\mathbf{D})]_q = D_q + \rho_q[(1-\psi)X_q(D_q) - R_q]$$

so that a steady-state equilibrium is $\bar{\mathbf{D}} = \{\bar{D}_q\}$ such that

$$\bar{\mathbf{D}} = T(\bar{\mathbf{D}}).$$

If T is a contraction mapping, then the contraction mapping theorem implies that it has a unique fixed point that can be computed by iteration. We conjecture, but as of this draft have not shown, that T satisfies Blackwell's sufficient conditions for a contraction, which would in turn be sufficient to finish proving the claim. □

B.2.3 ... of the full dynamic competitive equilibrium

Notice an important feature of the proof of Proposition B.4: the operator T is the law of motion for depth in Equation (9). We have shown that, by iterating this operator forward and solving for the unique trade equilibrium at each step, we are sure to converge to a unique steady state. The main result is then simply a corollary of the three preceding propositions.

Corollary B.1. *Given any initial vector of groundwater depths, $\{D_{q0}\}$, a unique competitive equilibrium exists.*

B.2.4 The outside sector as residual claimant

The aim of this section is to demonstrate that we demoted the outside sector from the equilibrium definition only out of convenience; no bugaboos are hiding there.

We've done some of the work already in previous appendices. As we showed in Appendix B.1.1, because the utility function in (1) is quasilinear, the representative consumer in country i spends her first ζ_i of income on agricultural output. All residual income is either (i) spent on the outside good or (ii) taxed away lump-sum.

In the main text we assumed conditions such that the outside good is always consumed and produced in every country. Moreover, we assumed that it is homogeneous and freely traded, so the law of one price holds (and we normalized that price to one). Accordingly, bilateral flows of the outside good are indeterminate; countries are simply net importers or net exporters from a unified global market for the outside good. In Appendix B.2.1, we stated the condition for that market to clear when there are no taxes or subsidies (namely, $z_o(\tilde{\mathbf{p}}) = 0$). That condition is guaranteed to be satisfied in equilibrium by Walras' law.

It just remains to check that reintroducing the distortions does no harm. The market clearing condition becomes

$$\sum_i (Y_{it} - T_{it} - \zeta_i) = \sum_i \sum_{f \in \mathcal{F}_i} h^f A_i^o \left(\pi_t^{fo}\right)^{\frac{\theta-1}{\theta}}$$

where $Y_{it} = Q_{it}^o + \sum_k \tau_{it}^k p_{it}^k Q_{it}^k$ is total income and T_{it} is the lump-sum tax. But notice that budget balance would require

$$\begin{aligned} T_{it} &= \sum_k (\tau_{it}^k - 1) p_{it}^k Q_{it}^k \\ &= \sum_k \tau_{it}^k p_{it}^k Q_{it}^k - \sum_k p_{it}^k Q_{it}^k, \end{aligned}$$

so market clearing with government budgets balanced is just

$$\sum_i \left(Q_{it}^o + \sum_k p_{it}^k Q_{it}^k - \zeta_i \right) = \sum_i \sum_{f \in \mathcal{F}_i} h^f A_i^o \left(\pi_t^{fo}\right)^{\frac{\theta-1}{\theta}}$$

which takes the exact same form as it did before we reintroduced taxes and is therefore satisfied by Walras' law.

C Figures Appendix

Figure C.1: Global distribution of water markets

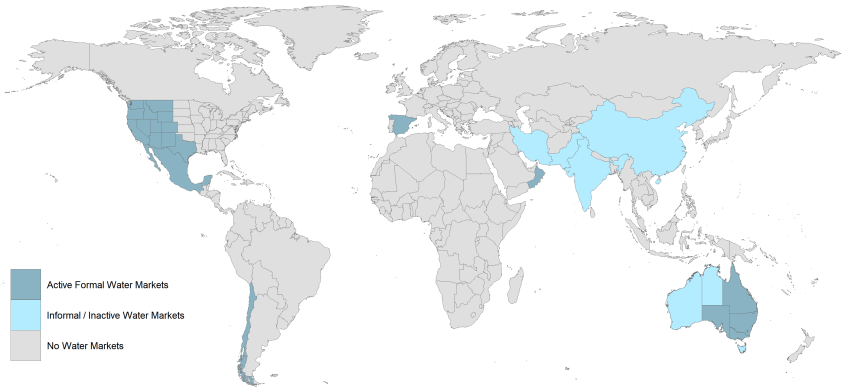

Notes: This map shows the global distribution of water markets, as uncovered through an extensive literature review. Subnational regions and countries shaded in dark blue have established formal water markets in which extraction rights can be traded. Regions and countries in light blue indicate locations where informal or inactive water markets are present. Grey areas indicate no formal or informal water markets exist. Water markets rarely extend across the entire regions shaded, but precise extents of these markets are impossible to determine with available data.

Figure C.2: Cropped Area Fraction by Decile of Water Variables

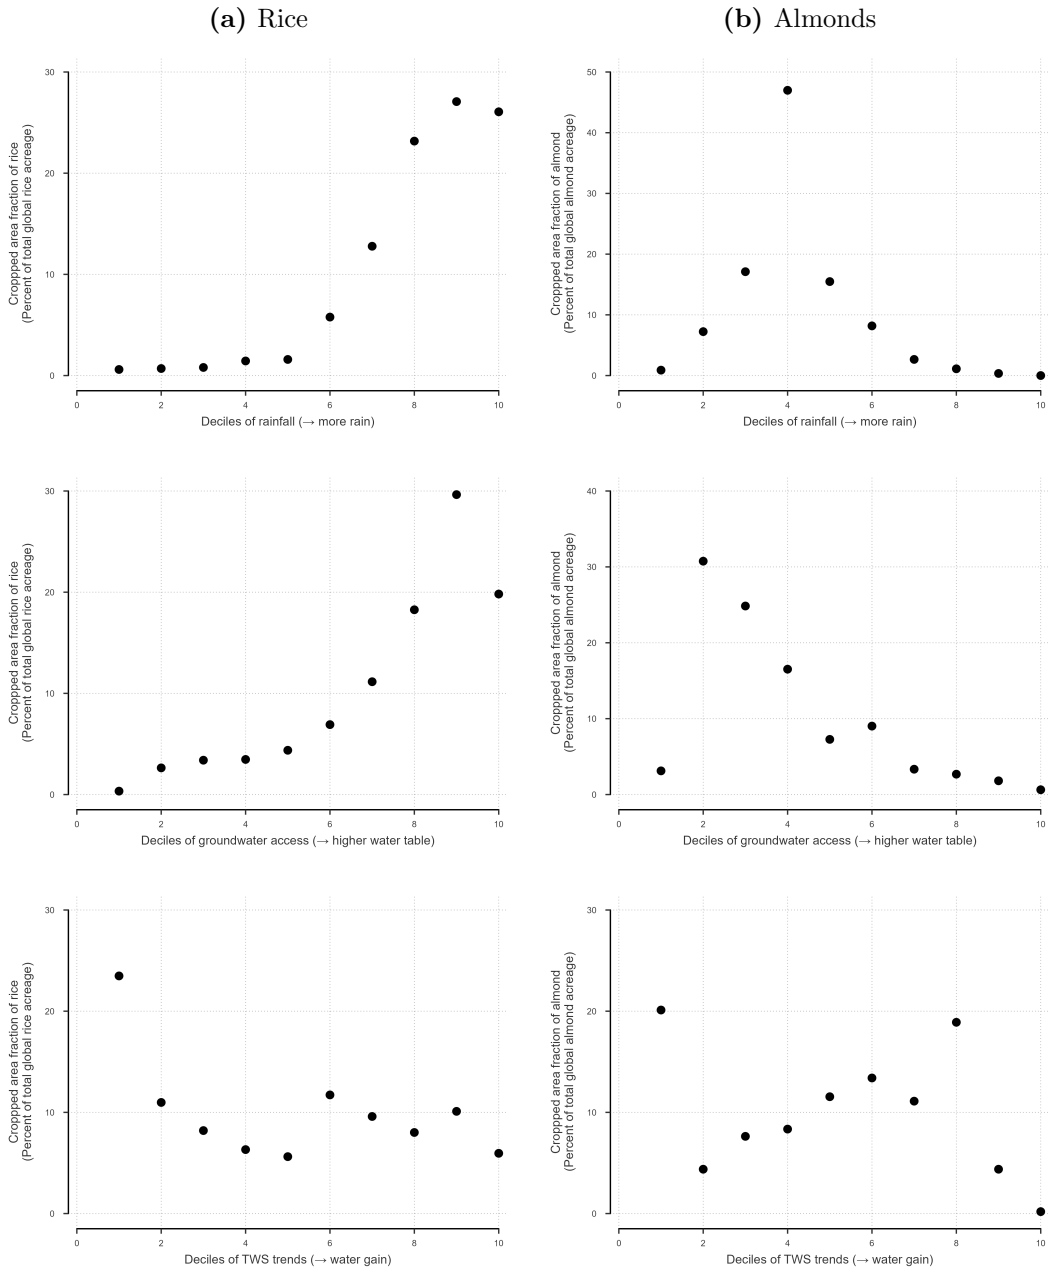

Notes: The three graphs in Panel (a) show the cropped area fraction of rice as a percent of total global rice acreage by decile of precipitation, groundwater table depth, and change in total water storage. Similarly, the cropped area fractions of almonds by decile of water variables are shown in the three graphs of Panel (b). Monfreda, Ramankutty and Foley (2008) compiled the crop-level agricultural use measure; the sources of the water variable data are discussed in Section A.