Measuring Water Use in Irrigated Agriculture with Satellites

Jiaxin Yin[1], Leslie Sanchez[1], Conner Harwood[2], Eric Edwards[1,2], Bryan Leonard[2,3], Will Rafey[4], Andrew Ayres[5], and Katherine Wright[6]

[1]Department of Agricultural and Resource Economics, University of California, Davis
[2]Property and Environment Research Center
[3]University of Wyoming
[4]University of California, Los Angeles and NBER
[5]University of Nevada
[6]Hillsdale College

August 18, 2026

Abstract

Satellite-based estimates of evapotranspiration (ET) are increasingly used to measure agricultural water use. Accurately measuring water use is essential for understanding the value of water and surprisingly few jurisdictions measure field level water use. Currently, relatively little is known about how effectively satellite ET can be used to characterize irrigator behavior. We systematically compare satellite-based estimates with traditional crop coefficient estimates from agronomy. Using field-level data from California and Utah, we demonstrate potential approaches for calculating the ET of applied water, a necessary modification of the raw ET data often used directly by economics researchers. We show that satellite ET estimates can differ substantially from expected agronomic behavior, but the magnitude and direction of these differences vary across crops and settings. In Utah, monthly comparisons show that much of the divergence reflects the timing of estimated water use. Agronomic crop coefficient methods generally impose shorter, more concentrated growing seasons, while satellite ET approaches lead to flatter seasonal profiles and water use extending later into the season. We recover the implied satellite-based crop coefficients, allowing these differences to be expressed directly within the framework used by crop scientists. We suggest that researchers consider the extent to which satellite ET models accurately characterize irrigator behavior and use multiple methods to demonstrate the robustness of policy evaluations or behavioral implications across approaches. We then compare three methods for separating irrigation-driven ET from naturally occurring ET: a structured water-balance model, a machine-learning model, and a parsimonious fixed-effects approach. Despite substantial differences in their structure and computational requirements, the methods generate similar distributions and seasonal patterns of applied ET in Utah. These results suggest that relatively simple statistical methods can recover much of the information provided by computationally intensive alternatives.

1 Introduction

Recent advances in remote sensing have made it possible to estimate crop water use at the field scale using satellite observations. The emergence of these satellite-based evapotranspiration (ET) products has been a notable advance for agricultural and natural resource economists. Researchers increasingly use satellite ET to estimate irrigation water demand (Barker et al., 2025), evaluate groundwater management policies (Lachhab and Bruno, 2026), measure the effectiveness of water conservation programs(Wright et al., 2026; Picciotto, 2025), estimate the value of water in agricultural production (Edwards and Sanchez 2026), and study water markets and environmental flows (Safari et al.).

Agencies and policymakers have also begun to incorporate ET products into water accounting and regulatory decisions (Volk et al., 2024; Pearson et al., 2024). Even in regions where water diversions are directly measured, which is surprisingly uncommon, calculating how much water is ultimately consumed by crops requires additional assumptions. As a result, satellite-based estimates of ET have the potential to improve our understanding of agricultural water use.

Despite the widespread use of ET products, relatively little attention has been given to the measurement properties of satellite ET. As a result, researchers often choose an arbitrary ET product without a clear understanding of what different measures of ET mean, when satellite-based estimates provide accurate measures of consumptive water use, when they diverge from traditional approaches, and how those differences may affect statistical inference.

The current ambiguity about the application of ET is apparent in the wide variety of methods used both to estimate ET itself and to convert those estimates to consumptive use—the portion of water withdrawn or diverted that is not returned to the water system and therefore not available for subsequent use. Without a better understanding of the strengths and limitations of satellite ET approaches, it is difficult to know when satellite ET provides a reliable measure of agricultural water use and when it may instead introduce systematic bias.

Our objective in this paper is to elucidate when and how satellite-derived ET can be used to study irrigator behavior. Doing so requires understanding the approaches for estimating ET, the sources of error that arise in different approaches to estimating crop water use. We address this question by comparing satellite-based ET estimates with estimates derived from crop coefficients (Crop ET), the standard approach used in crop science, irrigation scheduling, and many water management applications. Satellite methods are subject to measurement error arising from the remote sensing process itself as well as the algorithms by which raw data are converted into ET. In contrast, crop ET methods rely on observed crop characteristics and experimentally derived crop coefficients, so the primary concern becomes specification error. Crop coefficient methods require assumptions about crop water requirements, growing seasons, and management practices that may not accurately reflect conditions in a particular field or year.

We examine the ET of irrigated agriculture in two distinct settings: California and Utah. California represents a relatively high-value, technologically advanced agricultural system characterized by standardized growing seasons, extensive irrigation information, and a diverse set of

specialty crops. Utah is dominated by lower-value forage crops, particularly alfalfa; irrigation systems are generally less technologically intensive; growing seasons vary substantially with elevation; and centralized irrigation scheduling information is more limited. Comparing the estimates of satellite and crop coefficient approaches in these settings provides insight into the robustness of each approach and the conditions under which they can be used to measure agricultural water use.

2 Constructing Measures of Crop Water Use

2.1 Measuring Evapotranspiration for Economic Analysis

Evapotranspiration (ET) is the sum of two physical processes that transfer water from the ground to the atmosphere: evaporation of water from soil and plant surfaces and transpiration of water that plants draw up through their roots and release through their leaves. For an agricultural field, ET is effectively the water that leaves the system as vapor rather than running off, percolating below the root zone, or remaining in storage. ET is closely linked to plant growth and water consumption, so the amount of ET a field generates over a growing season is a close proxy for the water the crop actually uses. This makes ET a uniquely useful measure for economists studying agricultural water use.

Economists are interested in directly measuring consumptive water use as a key input into the plant growth process as well as a key aspect of water right allocation and transfer. Existing data on water use often relies on self-reported or administrative withdrawal records, which can be inaccurate, inconsistently reported, or unavailable. This is an important barrier to policy evaluation. A growing set of public programs, from groundwater trading arrangements to conservation grants and fallowing payments, are designed explicitly to reduce consumptive use but it is impossible to assess their performance without consumptive use data. Developing an empirical approach to estimating consumptive water use at the field level offers a promising way to fill this gap by generating policy-relevant evidence on water conservation, even in settings without direct monitoring infrastructure.

Historically, the main obstacle to using ET in economic research was measurement. On-the-ground ET measurement requires specialized instrumentation, such as flux towers, that is expensive to install and maintain. This made field-level, basin-wide, or multi-year ET data infeasible for a lot of applications. Remote sensing has changed this by observing the surface characteristics that determine the energy balance driving evapotranspiration, like surface temperature and vegetation indices. A family of models has been developed to convert these observations into ET estimates at the pixel level, at weekly-to-monthly frequency, across large geographic extents and back through history. This makes it possible to construct the type of field-by-month (or year), multi-year panel of water consumption that is increasingly enabling a broader research agenda connecting water consumption to agricultural production, land use, and policy.

2.2 OpenET and Remote-Sensed ET Models

OpenET is a platform that computes and disseminates field-level ET estimates using satellite data from the Landsat program, combined with weather data and land surface information. OpenET provides estimates from six independently developed ET models, each of which uses a different approach to convert satellite-observed surface conditions into an ET estimate [1]. All models, except SIMS, are based on surface energy balance (SEB) approaches that infer ET from the energy required to convert liquid water to vapor. Because the underlying physical logic is similar, the SEB-based OpenET models tend to be highly correlated but differ in their specific parameterizations and thus in their level and variance, particularly at very high or low levels of ET.

The ET output OpenET reports is *total* ET, or often referred to as *actual* ET, ET^a: total water vapor loss from a field, regardless of the water's source. Total ET, whether measured by satellite or otherwise, is not the same as irrigation use. A field that receives no irrigation at all can still have high ET^a in a wet year, since precipitation and stored soil moisture both support plant growth and evaporation independent of any water applied by an irrigator. To recover the water applied through irrigation, ET^a must be decomposed into the portion satisfied by non-irrigation sources, i.e., soil moisture and effective precipitation, and the residual, ET of applied water, $ET^{applied}$, which is the consumptive use of irrigation water.

Three of the four methods we consider, Jacobs, Statistical Subtraction, and Boser, start from the OpenET-derived ET^a and differ in how they separate applied from natural ET, as we describe below. The fourth, crop ET, is constructed independently of OpenET altogether. This approach builds up total ET from first principles, using reference ET (ET^o) and crop coefficients rather than satellite-observed surface conditions.

2.3 Constructing ET of Applied Water

To understand irrigator behavior, such as estimating how much water a policy conserves or how growers respond to water scarcity, simply grabbing an estimate of ET^a from OpenET will provide estimated levels that will be difficult to interpret. We subtract natural ET, ET^n, from ET^a to isolate the component of water consumption that is actually responsive to irrigation choices.

Figure 1 summarizes the four methods we consider for making this separation. All four target the same underlying object, $ET^{applied} = \text{Total ET} - \text{Natural ET}$, but they differ in their data inputs, in how they construct natural ET, and in whether that construction is mechanical, reduced-form, or machine-learning based: Jacobs relies on an explicit accounting of effective precipitation and carryover soil moisture; Statistical Subtraction estimates natural ET via a regression fit on fallowed fields; Boser predicts natural ET on fallowed pixels using machine learning over field characteristics; and crop ET sidesteps the decomposition entirely by building applied ET directly from crop coefficients and reference ET, net of effective precipitation.

We elaborate on the Jacobs and crop ET constructions, our primary and benchmark methods,

[1] eeMETRIC, geeSEBAL, PT-JPL, SIMS, SSEBop, and an ensemble that takes the weighted average of them

Figure 1: Workflow of the four ET of applied water construction methods

	Model-Based (Crop ET)	Statistical Subtraction	Boser	Jacobs
Step 1: Construct Total ET				
Data Inputs	Crop choice Crop coefficients (Kc) Reference ET Crop-specific irrigation schedule	Satellite-derived evapotranspiration: ALEXI/DisALEXI, eeMETRIC, geeSEBAL, PT-JPL, SIMS, SSEBop, Ensemble		
Construction method	Reference ET x Crop coefficient	Observed directly from satellite data		
Step 2: Construct Natural ET				
Land use class	–	Fallowed status	Fallowed status	Crop choice
Water availability	Precipitation	Precipitation	Precipitation	Precipitation, available water storage
Field characteristics	–	–	Elevation, slope, aspect, TWI Longitude, latitude Potential ET Soil characteristics	–
Construction method	Effective precipitation	Predicted ET on fallowed fields (via linear regression)	Predicted ET on fallowed pixels (via ML)	Effective precipitation + carry-over soil moisture
Step 3: Construct ET of Applied Water				
$ET^{applied} = Total\ ET - Natural\ ET$				
Methodological approach	Mechanical	Reduced form	Machine learning	Mechanical

below, and defer the details of Statistical Subtraction and Boser, which we treat as robustness checks.

2.3.1 Converting Satellite ET Measurements into Irrigated Crop Use

Satellite ET data, such as the OpenET products, report total water lost to the atmosphere through evaporation and transpiration over a given period at the pixel level. We refer to this as total ET, ET^a. Most OpenET models rely on a surface energy balance (SEB) approach, which infers ET from the energy required to convert liquid water to vapor using satellite-derived surface characteristics and weather variables, regardless of whether that water originates from precipitation, background soil moisture, or irrigation.

We estimate $ET^{applied}$ following the Jacobs method, developed by Jacobs Engineering for the Colorado River Authority of Utah's Demand Management Pilot Program. Jacobs arrives at applied ET by explicitly modeling the other sources of crop water demand: soil moisture and effective precipitation, and subtracting them from total ET:

$$ET^{applied}_{imy} = \max\{0,\ ET^{a}_{imy} - SM_{imy} - P^{eff}_{imy}\} \tag{1}$$

where SM_{imy} is available soil moisture and P^{eff}_{imy} is effective precipitation for field i in month m and year y. For non-irrigated fields, $ET^{applied} = 0$. This makes Jacobs the more structured, engineering-driven of the satellite ET approaches we consider, in contrast to the machine-learning-

based and reduced-form alternatives discussed in Section 5.

We obtain ET estimates from OpenET's eeMETRIC model for the contiguous United States via Google Earth Engine, at monthly frequency, averaged to the field level to construct ET^a.

Effective precipitation is calculated following USDA (1970):

$$P^{eff}_{imy} = SF_i(0.70917P^{0.82416}_{imy} - 0.11556)(10^{0.02426ET^a_{it}}) \tag{2}$$

where SF_i, P_{imy}, and ET^a_{imy} are field i's soil water storage factor, monthly precipitation, and monthly ET, respectively. Precipitation is drawn from Oregon State University's PRISM dataset at 800 m resolution, aggregated to monthly spatial means. Soil characteristics are drawn from USDA's SSURGO database, from which we calculate the soil water storage factor as

$$SF_i = 0.531747 + 0.295164D_i - 0.057697D^2_i + 0.003804D^3_i \tag{3}$$

where D_i, usable water storage, is taken as 40% of field i's available water storage for a soil depth of 0 to 59 inches.

Soil moisture in Equation 1 begins each irrigation season (April-October) as a carryover value from the preceding winter (November-March) and is sequentially depleted over the irrigation season. We calculate carryover soil moisture for each field-year following Hill et al. (1989):

$$SM^{co}_{iy} = \min\{0.67(P^{win}_{iy} - 1.25ET^{win}_{iy}),\ 0.75 \times RZ_{iy} \times AWC_i\} \tag{4}$$

where P^{win}_{iy} and ET^{win}_{iy} are field i's total winter precipitation and ET, respectively, and RZ_{iy} is the rooting zone depth of the majority crop grown on field i in year y, taken as the mean of the maximum range in Allen et al. (1998) as recommended by Barker et al. (2025), with the Jacobs Method documentation used for crops missing from that source. Soil available water capacity, AWC_i, is available water storage divided by 59.

We track the month-to-month drawdown of soil moisture with a sequential water balance model. At the start of each April, SM^{co}_{iy} initializes available soil moisture, and Equation 1 is applied to estimate April consumptive use. Soil moisture is then depleted by $\min\{SM_{imy},\ \max\{0,\ ET^a_{imy} - P^{eff}_{imy}\}\}$, the amount used to satisfy crop water demand not already met by effective precipitation, and the resulting carryover initializes May. This process repeats through October; once soil moisture reaches zero, applied ET is calculated as total ET in excess of effective precipitation alone.

2.3.2 Constructing Crop ET Using Crop Coefficients and Reference ET

Crop ET models ET based on crop growth stage and local weather conditions. This is the standard approach in crop science and irrigation scheduling. The construction of crop ET combines two ingredients: reference ET, ET^o, and crop coefficients, K_c. Reference ET measures the evaporative demand of the atmosphere under standardized conditions: how much water a well-watered grass surface would consume given local weather. The crop coefficient scales this reference demand to

the specific crop and the growth stage. We assign a crop-specific coefficient for pixel i growing crop c at date t, $K_{c,it}$, based on the irrigation schedule and combine it with daily reference ET to obtain the crop water demand, which we define as total crop ET:

$$ET^c_{c,it} = ET^o_{it} \times K_{c,it}. \tag{5}$$

Note that the crop coefficient, K_c, is not a fixed number but a curve that traces a crop's water demand across the growing season, rising from the crop initiation, peaking near full cover, and falling as the crop goes to termination. K_c is assigned at each point in time based on where the crop falls at each growing stage, so the accuracy of ET^c depends both on the shape of the coefficient curve itself and on correctly dating the time of initiation, full cover, and termination for a given crop, region, and year.

Daily values are integrated to obtain monthly water demand, and the pixel-level estimates are aggregated to the field level. We then subtract effective precipitation from the field-month estimate. Repeating this across all pixels and months yields a panel of crop ET at the field-month level.

2.4 Data Construction

We examine irrigated agriculture in two settings: Utah and California. We construct the applied ET panel for two methods at the field-month level separately for each setting.

2.4.1 Utah

Utah agriculture relies heavily on lower-technology irrigation systems, grows relatively low-value crops (primarily alfalfa), has variable growing seasons due to high elevation, and has more limited centralized production information.

The Utah sample spans seven years (2017-2023) and covers about 116,000 fields totaling 1.95 million acres in the most recent year. Agriculture in Utah is dominated by hay and turf crops, primarily alfalfa (33% of total agricultural acreage) and grass hay (13%), followed by pasture (22%) and fallow/idle land (18%). Grain and seed crops (mostly winter wheat and barley, 9%) and field crops (mostly corn, 5%) account for most of the remainder, with orchard and small-fruit acreage negligible (<1%). Fields are irrigated primarily by flood (43%) and sprinkler (35%) methods, with drip irrigation rare (<1%), reflecting Utah's reliance on less efficient gravity-fed canal delivery rather than the highly efficient, low-flow micro-irrigation systems common among California's specialty crops. Due to frosts, Utah's agricultural season is usually constrained to April-October, with variable initiation and termination dates for different crops across regions. total ET estimates (OpenET, eeMETRIC) across all Utah fields range from 0.55 ft/yr (5th percentile) to 3.1 ft/yr (95th percentile), with a median of 2.1 ft/yr.

Utah field boundaries are drawn from the Water-Related Land Use (WRLU) dataset in the Utah State Geographic Information Database (UCGID), anchored to 2023 field boundaries to retain only

stable fields over the 2017-2023 study period.

Crop categories and types are identified using the Cropland Data Layer (CDL). Reference ET and weather variables are drawn from gridMET (University of Idaho Climate Research Group), providing daily data at 4-km resolution across the contiguous United States. Crop coefficient curves follow the engineering estimates in Hill et al. (2011). Because the frost period is variable in Utah, no planting dates are pre-specified. Instead, the irrigation schedules are constructed from three growing-timing controls: initiation, effective full cover (EFC), and termination Hill et al. (2011). We first compute daily Hargreaves reference ET, ET^H, using gridMET weather data and each pixel's latitude, longitude, and day of year, following FAO-56 specification (Allen et al., 1998):

$$ET^H = 0.0023(T_{mean} + 17.8)(T_{max} - T_{min})^{0.5} R_a, \tag{6}$$

where T_{mean}, T_{max}, and T_{min} are daily mean, maximum, and minimum temperature, R_a is extraterrestrial radiation, and 0.0023 and 17.8 are empirically determined scaling constants. Growth initiation is set when accumulated ET^H since January 1 crosses a crop-specific threshold and temperature conditions permit growth. EFC, the point of maximum water consumption, is set either a fixed number of days or a fixed number of growing degree days after initiation, depending on the crop's coefficient curve. Termination is set by days from EFC to harvest, cumulative growing degree days, or a crop-specific killing-frost temperature.

2.4.2 California

California agriculture, on the other hand, is relatively high value and technologically advanced, with standardized growing seasons, centralized irrigation information, and is largely dominated by high-value crops.

The California sample spans six years (2018-2023) and covers roughly 285,000 fields totaling 6.53 million acres in the most recent year. Perennial and permanent crops dominate the acreage: deciduous fruit and nut orchards alone account for 30% of the sample, followed by pasture (15%), field crops (9%), vineyards (9%), truck/nursery/berry crops (7%), rice (7%), grain and hay crops (6%), and citrus/subtropical crops (5%). Fallow land makes up 2.8% of total agricultural acreage, and a further 7.9% is classified as idle. While surface irrigation (flood and furrow) still dominates alfalfa and pasture, drip and micro-sprinkler are widely applied in the production of high-value specialty crops, such as almond, grape, and citrus. California spans several distinct hydrological regions, and the growing season varies by region: year-round in the Colorado River basin, February-October along the south coast, and April-October elsewhere in the state. Total ET estimates (OpenET, eeMETRIC) across all fields range from 0.85 ft/yr (5th percentile) to 3.7 ft/yr (95th percentile), with a median of 2.1 ft/yr. Table 1 gives crop-specific ET ranges for the major crops grown in both California and Utah.

Field boundaries are drawn from the California Department of Water Resources (CDWR) i15 crop mapping dataset, anchored to 2022 boundaries to retain fields stable in location and size

over the 2018-2023 study period. Crop categories and types are identified through LandIQ, cross-referenced against the Cropland Data Layer (CDL).

Crop coefficients and irrigation schedules are assigned by region and crop using a pre-specified planting date and coefficient curve. These come from the engineering estimates specified in Leaflet 21427 by Cooperative Extension, University of California, Division of Agriculture and Natural Resources.

Table 1: Study area and data summary

	California	**Utah**
Study period	2018-2023	2017-2023
Unique fields	285072	116468
Total acreage (1M acres)[a]	6.53	1.95
Fallow reference share (%)	2.8	18.1
Total ET, 5th / 50th / 95th pctile (ft/yr)	0.85 / 2.10 / 3.69	0.55 / 2.06 / 3.10

California

Rank	Crop / class	Share	ET range[b] p5-p95	(med.)
1	Deciduous fruits and nuts	29.9%	1.15-3.98	(2.67)
2	Pasture	15.3%	0.91-4.14	(2.42)
3	Field crops	9.2%	1.95-3.64	(2.84)
4	Vineyards	9.1%	0.76-3.38	(1.62)
5	Idle	7.9%	0.47-2.66	(1.39)
6	Truck, nursery, and berry crops	6.9%	1.13-3.08	(2.07)
7	Rice	6.8%	2.67-3.82	(3.27)
8	Grain and hay crops	6.4%	0.72-2.85	(1.51)
9	Citrus and subtropical	5.2%	1.30-3.27	(2.20)
10	Unclassified fallow	2.8%	0.52-2.68	(1.39)

Utah

Rank	Crop / class	Share	ET range[b] p5-p95	(med.)
1	Alfalfa	33.0%	1.10-3.22	(2.34)
2	Pasture	21.6%	0.43-2.99	(1.83)
3	Fallow/Idle	18.1%	0.26-2.29	(1.06)
4	Grass Hay	13.3%	1.07-3.15	(2.30)
5	Winter Wheat	5.9%	0.71-2.66	(1.73)
6	Corn	4.5%	1.48-2.82	(2.20)
7	Triticale	0.7%	0.73-2.94	(1.84)
8	Safflower	0.7%	0.81-2.02	(1.39)
9	Barley	0.6%	1.09-2.62	(1.93)
10	Oats	0.3%	0.94-2.92	(1.98)

[a] Acreage shares are a single-year snapshot (2023, the latest year in each panel), not a multi-year average.

[b] Field-year total ET (OpenET, eeMETRIC model), 5th-95th percentile and median, in ft/yr.

3 Comparing Satellite ET and Crop ET

3.1 Comparison of Annual Applied ET Methods

We begin with a simple comparison of the mean annual applied ET estimates using the Crop ET and satellite ET approaches. Figure 3.1 compares estimates using the satellite-based Jacobs method and the Crop ET method across major crop groups in California (panel a) and Utah (panel b). The figure shows both the estimates of the ET of applied water as well as the underlying total ET from the satellite approach and the total crop ET ($ET^o \times K_c$) from the Crop ET approach. The two approaches often produce estimates of similar magnitude, but the degree of agreement varies between crops and states. The level of applied ET differs between crops, with field crops, fruits, and nuts in California using the most water while alfalfa uses the most in Utah.

Neither method is systematically higher across all crops in either state. In California, the methods arrive at similar estimates for field crops and rice. Crop ET is higher for citrus and subtropical crops, pasture, and deciduous fruits and nuts; satellite estimates are higher for grain and hay crops and truck, nursery, and berry crops. In Utah, Crop ET estimates are higher for alfalfa while satellite ET estimates are higher for other crop groups.

Although our ultimate goal is finding the ET of applied water, all the differences in estimated water use come before converting total ET into applied ET. In California, where precipitation during the growing season is extremely low, there is little divergence between total ET and ET of applied water or between crop ET and total crop ET ($ET^o \times K_c$). In Utah, which has more precipitation during the growing season, the reduction due to effective precipitation and carryover soil moisture is larger.

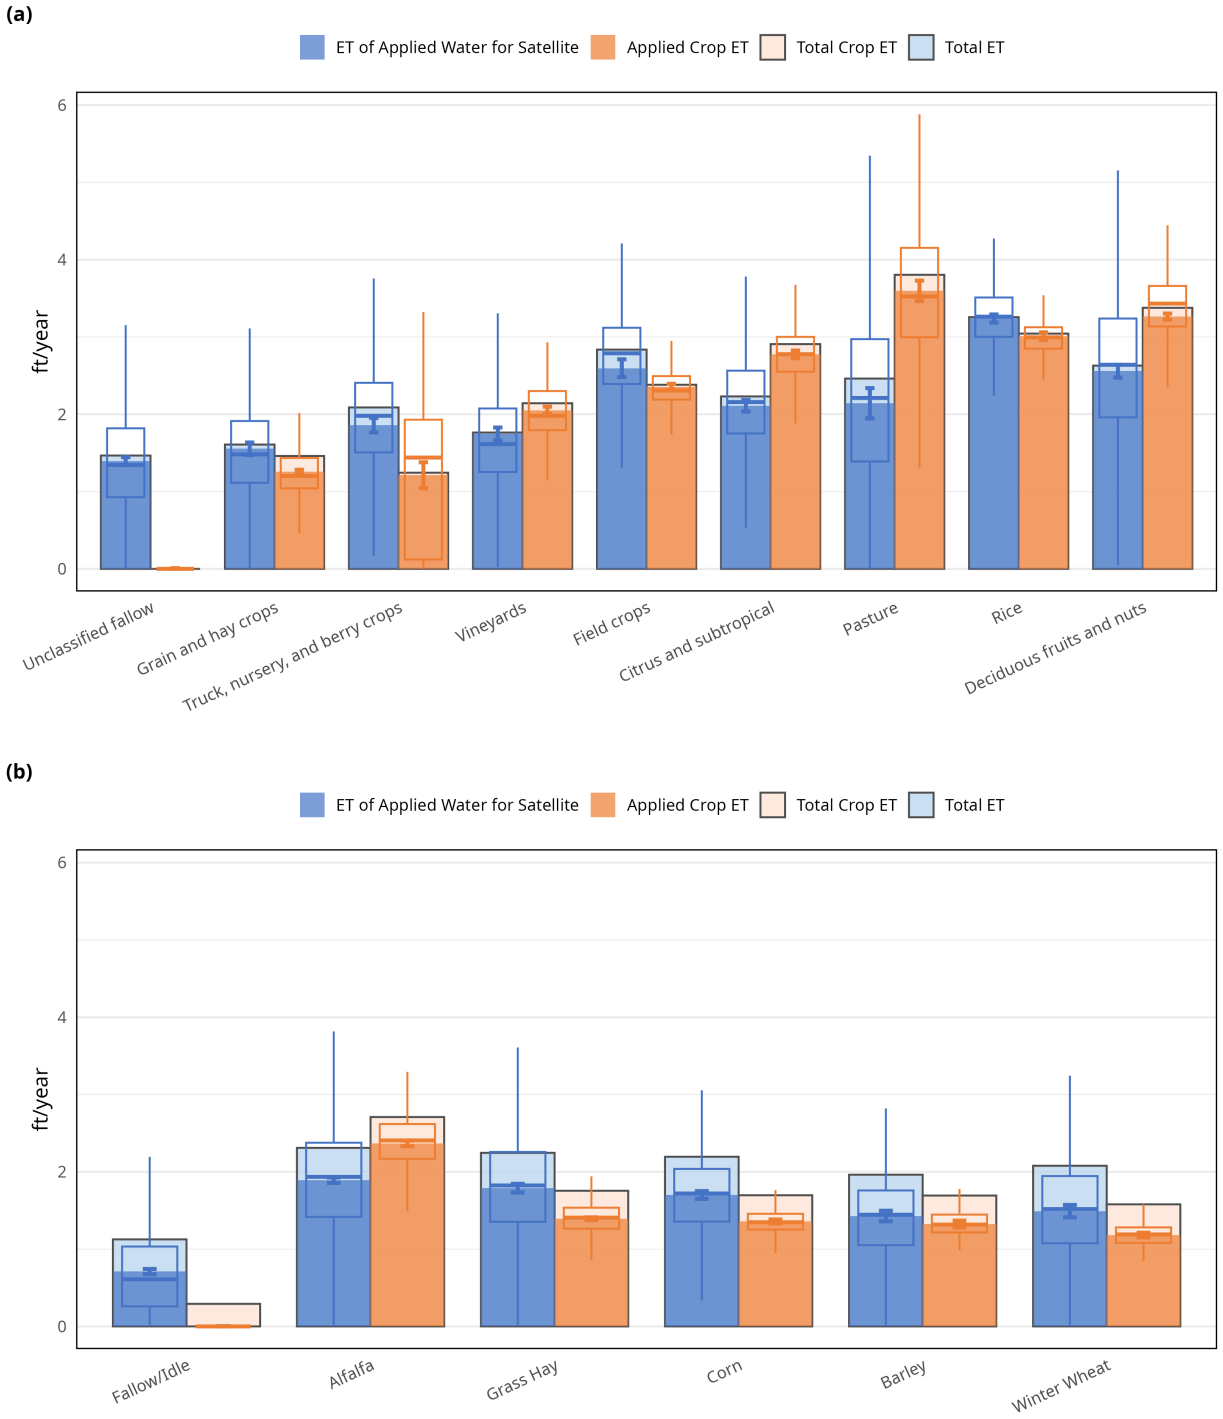
Figure 3.1 suggests that neither method produces uniformly higher estimates across settings. Instead, the differences are crop- and location-specific, consistent with differences in the measurement and specification assumptions underlying the two approaches. Fallow/idle land provides an important benchmark: the crop ET method, by construction, assigns no applied ET to these fields, while the satellite-based method detects positive ET attributable to water beyond modeled precipitation and soil moisture. We return to the question of whether fallowed fields could still have positive ET of applied water in section 5.3

3.2 Seasonal Sources of Divergence

Annual differences can obscure an important source of disagreement between the two methods: the timing of estimated water use within the growing season. Figure 2 therefore decomposes the Utah estimates by month from April through October. The comparison reveals that differences in annual applied ET often arise not primarily from disagreement about peak water demand, but from disagreement about the beginning, duration, and end of the irrigation season.

Alfalfa provides the strongest case of agreement. Both methods generate a similar seasonal profile, with applied ET increasing through the spring, peaking in midsummer, and declining through September and October. Crop ET is somewhat higher early in the season, but the esti-

Figure 2: **Yearly Applied ET Estimates by Crop, Satellite-Based (Jacobs) Method vs. Crop ET Method. for (a) California and (b) Utah.** Colored bars show the crop-year mean applied ET under each method (blue = Jacobs, orange = Crop ET); boxplots overlay the interquartile range of the underlying field-year distribution, and error bars give 95% confidence intervals from a regression of applied ET on crop and year fixed effects, clustered at the HUC-12 watershed level. The paired background bars show the two methods' underlying ET components for reference: raw OpenET total ET (light blue) behind the ET of applied water (Jacobs) bars, and total crop ET, $ET^o \times K_c$ (light orange), behind the applied crop ET bars. Crops are California's primary DWR land-use classes in (a) and Utah's major crop groups in (b).

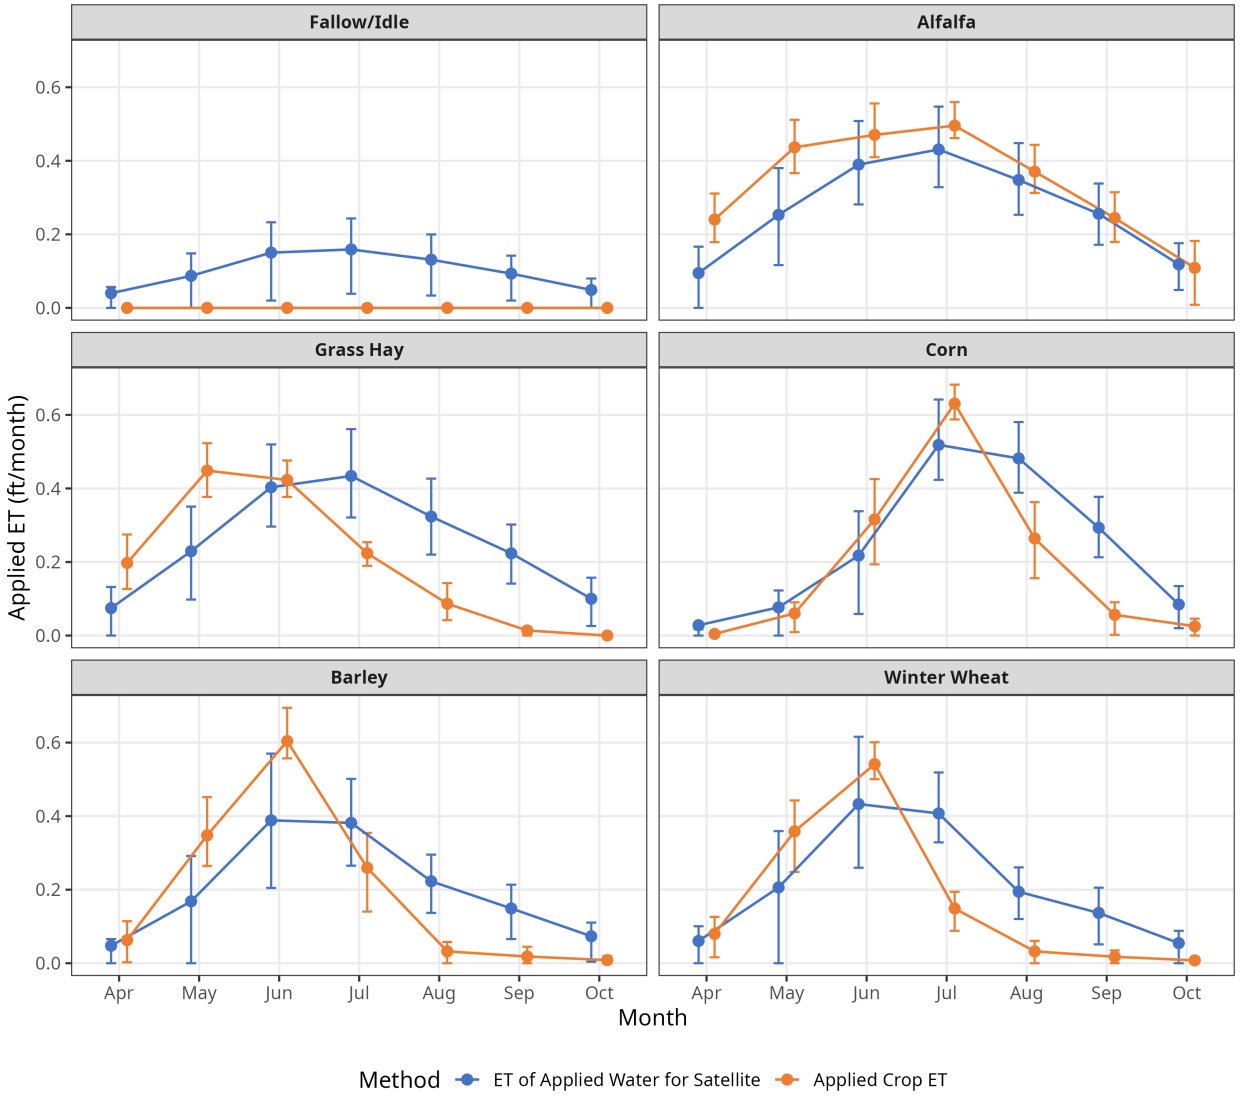


mates converge later in the summer. This close correspondence in both the level and timing of water use helps explain the similar annual estimates for alfalfa in Figure 1.

For the remaining crops, the seasonal profiles reveal considerably greater divergence. For grass hay, Crop ET implies a relatively concentrated irrigation season: estimated water use rises rapidly in May and June and then declines sharply, approaching zero by September. Jacobs instead implies a broader season, with lower estimated water use early in the spring but substantially more water use from July through October. A similar pattern appears for barley and winter wheat. Crop ET generates pronounced spring or early-summer peaks followed by rapid declines, while Jacobs produces flatter profiles with substantial applied ET persisting later into the season. The higher annual Jacobs estimates for these crops in Figure 1 therefore appear to be driven largely by this additional late-season water use.

Corn exhibits a related but somewhat different pattern. Both methods identify a strong mid-summer peak and generate similar estimates through July. Beginning in August, however, Crop ET declines rapidly, whereas the satellite estimate remains elevated through September and October. Thus, the annual difference for corn is again concentrated primarily in the assumed termination of the growing and irrigation season rather than in estimates of peak demand.

Figure 3: **Monthly applied ET by crop in Utah, April–October, comparing the Jacobs (satellite) and Crop ET (engineering) methods.** Points show the field-month mean applied ET for each method (blue = Jacobs, orange = Crop ET) and error bars span the interquartile range; panels facet by crop. The comparison highlights within-season timing differences between the two methods' estimated irrigation demand, in addition to the level differences summarized in Figure 1(b).



4 Reconciling Crop ET and Satellite ET

The seasonal differences in Figure 3 suggest that divergence between Crop ET and satellite ET estimates may arise from the crop coefficients used to construct the engineering estimates. Crop ET assumes that water demand follows a predetermined crop-specific curve, while satellite ET measures realized ET without imposing that seasonal profile. To make this comparison explicit, Figure 4 recovers satellite derived crop coefficient from the Jacobs estimates and compares it with the crop coefficients used to construct Crop ET. Because crop ET is defined as ($ET_c = K_c \times ET_o$), we estimate the satellite (K_c) by relating satellite-based crop ET to reference ET separately for each crop and month. The resulting coefficients describe the seasonal pattern of water use implied by the satellite data in the same units as the standard crop coefficients.

For alfalfa, the coefficients have broadly similar seasonal shapes, particularly during the middle and end of the growing season. The principal difference occurs early in the growing season. The standard coefficient begins substantially above the satellite coefficient in April and May, while the satellite coefficient rises more gradually. By midsummer the two converge and remain quite similar through October. This pattern is consistent with Figure 3, where Crop ET exceeds the satellite approach early in the alfalfa season but the two methods produce similar estimates later in the year.

For the other crops, the differences are much larger and systematic, although with the exception of grass the shapes of the coefficient curves are somewhat similar. The standard coefficients generally impose a pronounced peak followed by a rapid decline. Grass hay, for example, reaches a coefficient of roughly one in May before falling sharply after June. The satellite coefficient instead rises through early summer and then remains relatively flat through the remainder of the season. Barley shows a similar pattern: the standard coefficient peaks near one in June and falls close to zero by late summer, while the satellite coefficient peaks at a substantially lower level and remains around one-half through October. Winter wheat likewise has a high standard coefficient in May and June followed by a sharp decline, whereas the satellite coefficient declines much more gradually and remains positive through the fall.

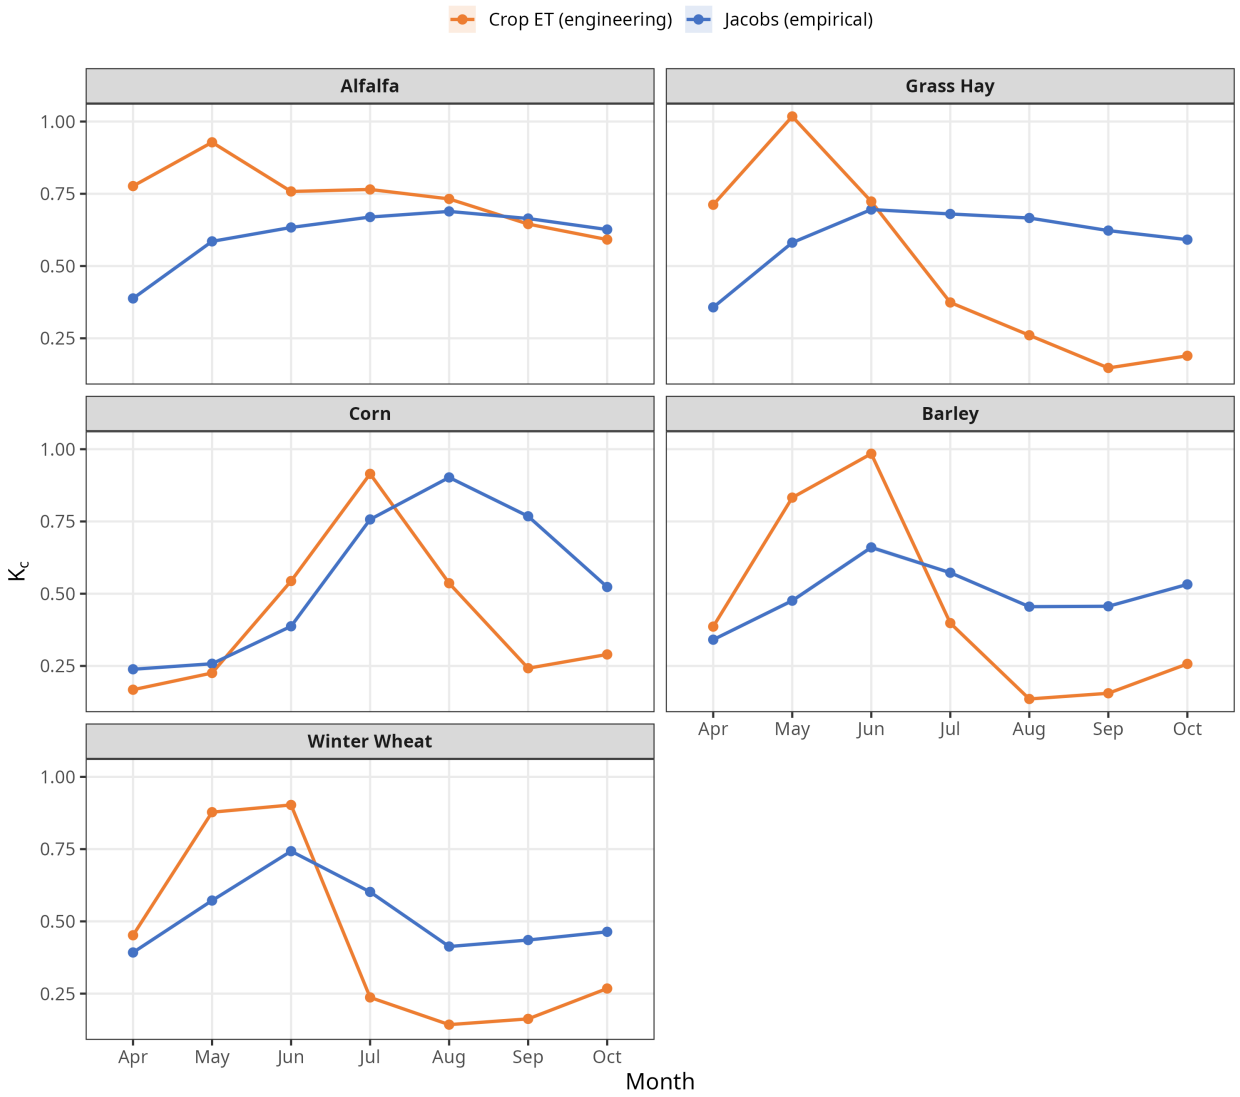
For corn, both approaches imply a clear midsummer peak. The standard coefficient rises sharply to a July maximum and then falls quickly, while the satellite coefficient peaks later and remains high into August and September. This again mirrors the monthly applied ET estimates in Figure 3 the two methods are relatively close around peak demand but the satellite coefficients imply substantially more late-season water use.

Taken together, Figure 4 indicates that the differences between Crop ET and satellite ET cannot generally be reconciled by a simple level adjustment to the crop coefficients. For most Utah crops, the satellite data imply a different seasonal shape. standard coefficients tend to concentrate water demand around a relatively short peak growing period, whereas the satellite coefficients are generally flatter and imply a longer period of water use. Alfalfa is the notable exception, with relatively close agreement once the crop reaches full cover.

These results point to the timing assumptions embedded in standard crop coefficients as an

important source of divergence between Crop ET models and satellite models. The satellite coefficients do not by themselves establish that the standard coefficients are incorrect. The satellite ET estimates suggest farmers in practice are less uniform and irrigate non-alfalfa crops more in the later season than expected (or potentially needed) by the crop ET model. Satellite ET methods are also measured with error, and differences are also consistent with less precise measurement and an aggregation individual behaviors into aggregate curves. What the exercise provides is a direct way to diagnose where the two approaches disagree and suggests the need for researchers to consider the importance of the choice of ET model in understanding irrigator behavior. It also suggests a potential role for satellite ET in updating crop coefficients to better reflect observed field-level water use, while preserving the underlying structure of the crop coefficient approach.

Figure 4: **Satellite versus empirical crop coefficients (Kc) by crop and month, Utah only.** Because crop ET is constructed as $ET_c = K_c \times ET_o$ and $ET_{applied} = ET_c -$ effective precipitation, regressing ET_c on ET_o through the origin, separately by crop-month cell, recovers an implied Kc. Applying this to the Crop ET method's applied ET recovers the engineering Kc used to construct it by definition (orange); applying it to the Jacobs satellite-based applied ET recovers an empirical Kc implied by observed field-level water use (blue). Lines show point estimates and shaded ribbons show 95% confidence intervals; panels facet by crop.



5 Alternative Satellite ET Methods

Method	Specification Error	Measurement Error
Crop ET	Total ET component	Limited in natural ET component
Boser	Limited	Total ET and natural ET components
Jacobs	Natural ET component	Total ET component
Subtraction	Limited controls relative to Boser	Total ET and natural ET components

5.1 Alternative Satellite ET Approaches

Jacobs is not the only way to construct applied ET from satellite ET. We consider two alternatives that solve the same underlying problem of inferring the portion of actual ET attributable to irrigation. The Boser method replaces Jacobs's explicit water balance with a machine learning model that predicts the natural-ET counterfactual directly, learned from fallow pixels rather than modeled from soil moisture and precipitation inputs. Statistical Subtraction method uses the same intuition of using fallow parcels to predict counterfactual natural ET for irrigated parcels, replacing Boser's gradient boosting model with a simple fixed-effects regression on fallow parcels, targeting the same fallow-to-irrigated prediction problem with a fraction of the computational cost.

5.1.1 Boser Method

The Boser method (Boser et al., 2024) recovers applied ET as the gap between observed actual ET and a counterfactual estimate of natural ET, ET^o, the ET that would have occurred had the irrigated pixel been fallow:

$$ET^{applied} = ET^a - ET^o \tag{7}$$

Because natural ET is unobserved on irrigated pixels, Boser approximates it using a machine learning model trained on fallow pixels, where $ET^a = ET^o$ by construction. The model learns the relationship between natural ET and a set of predictors, including topography, soil properties, potential ET, location, and time, and this relationship is then used to predict the counterfactual ET^o on irrigated pixels. Unlike Jacobs, which arrives at applied ET by explicitly accounting for soil moisture and precipitation, Boser predicts the natural-ET counterfactual directly and is accordingly less structured and more data-intensive.

We predict natural ET using a gradient boosting regressor trained on a filtered set of fallow pixels. Predictors include geographic coordinates (longitude, latitude), time (month, year), topographic features (elevation, slope, aspect, topographic wetness index), soil characteristics, and potential ET. We split the data into training (60%), validation (10%), and testing (30%) sets using spatial clustering: pixels are grouped into 2-km clusters, and each cluster is assigned entirely to one split, preventing nearby, similar pixels from appearing across splits. Hyperparameters are tuned via randomized search with 100 iterations of three-fold cross-validation on an unclustered subset.

Misclassification of fallow parcels can bias the natural ET model upward and thus understate applied ET. In Utah, fallow fields are identified through the CDL crop category, with two additional filters applied: parcels indicating subsurface irrigation are dropped, as these can elevate ET absent visible vegetation, and pixels with annual ET exceeding 80% of annual precipitation are dropped, on the assumption that natural ET sustained by precipitation alone is unlikely to exceed this share once runoff and percolation losses are accounted for. In California, fallow fields are identified as parcels classified as fallow (code X) in LandIQ and as Fallow/Idle in CDL, with pixels in the top 5% of the ET distribution, which imply abnormally high ET, dropped from the training set. The resulting cleaned fallow sets are used to train the natural ET prediction model in each setting.

5.1.2 Statistical Subtraction

Statistical Subtraction targets the same prediction problem as Boser: estimating natural ET on irrigated parcels using a model trained on fallow parcels, but it replaces Boser's gradient boosting model with a linear regression on year-month and region (county or HUC-12) fixed effects:

$$ET^{o}_{i,ym} = \alpha_{ym} + \gamma_{region(i)} + \varepsilon_{i,ym} \tag{8}$$

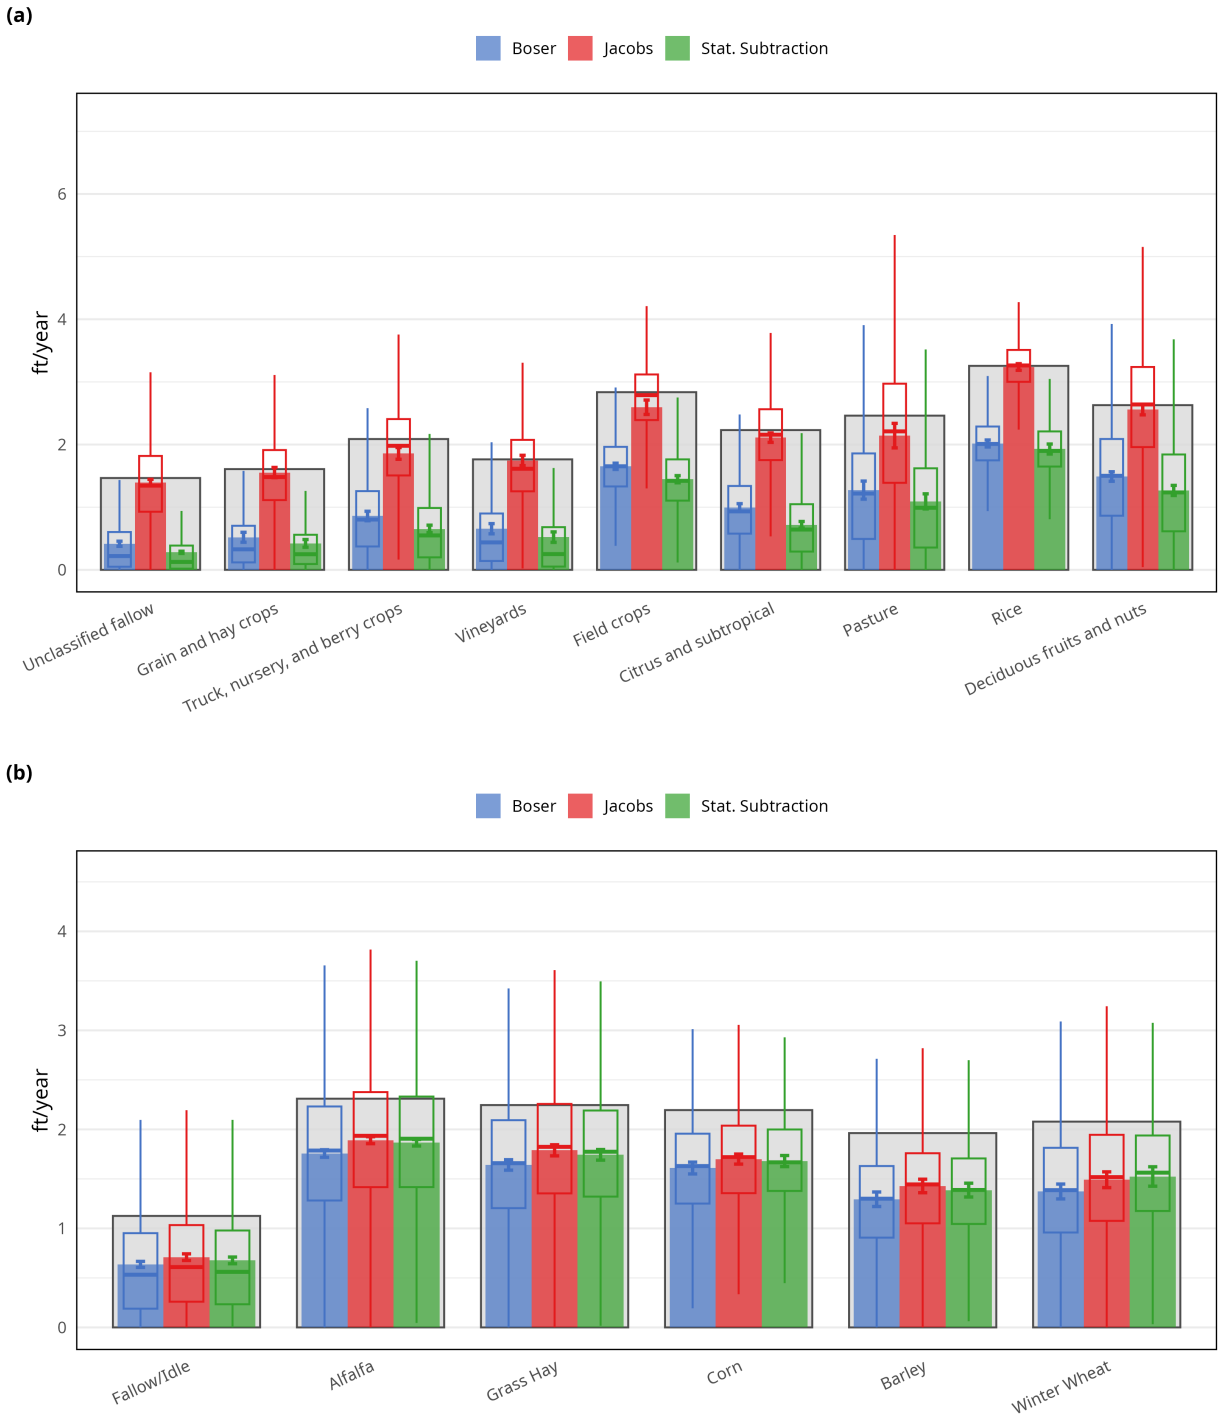
estimated on the sample of fallow parcels, where α_{ym} is a year-month fixed effect, $\gamma_{region(i)}$ is a county or HUC-12 fixed effect for the county or watershed containing field i, and $\varepsilon_{i,ym}$ is an idiosyncratic error term. Natural ET for irrigated parcels is then predicted using the estimated fixed effects corresponding to each parcel's location and year-month. Where Boser draws on a rich set of topographic, soil, and climate predictors, Statistical Subtraction relies only on coarse spatial and temporal averages, making it substantially simpler and less computationally intensive, closer in spirit to a reduced-form econometric specification than to a machine learning model.

5.2 Robustness Analysis Across Satellite ET Methods

The preceding analysis used the Jacobs method to convert satellite estimates of actual ET into applied ET. The differences between Jacobs and Crop ET could therefore reflect either the distinction between satellite and Crop ET estimates of crop water use or assumptions specific to the Jacobs method for separating irrigation from naturally occurring ET. Figures 5 and 6 evaluate this possibility by comparing Jacobs with the Boser and Statistical Subtraction methods.

Figure 5 compares annual applied ET across the three satellite methods. In Utah, the methods produce similar estimates. For alfalfa, grass hay, corn, barley, and winter wheat, the means and underlying distributions overlap substantially. The agreement is especially notable given the different approaches used to estimate natural ET:. The similarity of their estimates suggests that the annual differences between Crop ET and Jacobs documented in Section 3 are not primarily a consequence of the particular method Jacobs uses to remove naturally occurring ET.

Figure 5: **Robustness of the applied-ET method comparison to the choice of satellite ET product, Boser (blue), Jacobs (red), and statistical subtraction (green), for (a) California and (b) Utah.** Colored bars show crop-year means, boxplots show the interquartile range, and error bars show 95% confidence intervals from a crop- and year-fixed-effects regression clustered at the HUC-12 level; the grey background bar gives the mean raw OpenET actual ET for reference.



The California results show greater differences in levels. Jacobs generally produces higher applied ET than Boser and Statistical Subtraction, while Boser tends to produce the lowest estimates. The likely explanation for the divergence is that the Jacobs method is currently calibrated for Utah and underestimates natural ET.

Figure 6 shows that the agreement among satellite methods is quite strong when we consider within-season variation. Currently, these estimates are only available for Utah. The three approaches generate nearly identical seasonal profiles for each major crop. They identify the same timing of peak water use and the same persistence of water use later in the season. The broad late-season tails for grass hay, corn, barley, and winter wheat that distinguish satellite ET from the Crop ET estimates in Figure 3 appear under all three satellite methods.

Together, Figure 5 and 6 suggest an important distinction between uncertainty about the level and the temporal pattern of satellite-based applied ET. Alternative methods for removing natural ET can generate meaningful level differences, especially in California. In contrast, the cross-crop rankings and within-season patterns are comparatively robust because this variation largely stems from the total ET estimates.

5.3 Counterfactual Selection Check

Throughout both Boser and Statistical Subtraction approaches, fallow/idle fields serve as the "dry" counterfactual for irrigated fields. This relies on an important assumption that the fallow/idle fields have zero, or nearly zero, applied ET. If a fallow field sits adjacent to actively irrigated land, it could receive spillover water through runoff, seepage, or leakage. Any of which would inflate the reference baseline and bias the applied ET toward understating true crop water use. Figure 7 tests this directly. Each fallowed field is matched to its eight nearest irrigated fields in the same year, and boxplots show the fallow parcel's raw annual ET binned by distance to its nearest irrigated neighbor, for California (panel a) and Utah (panel b).

In both states, median fallow ET declines as distance to the nearest irrigated field increases, with the decline leveling off beyond roughly one kilometer. Fallow parcels sitting within 100 meters of irrigated land show visibly higher ET than parcels more than a kilometer away, consistent with some degree of spillover near the irrigated/fallow boundary. Beyond that distance, however, the relationship flattens rather than continuing to fall. Therefore, the spillover effect is geographically localized. In our paper, we pool all fallow fields together as a single counterfactual reference class, lowering bias into the applied ET estimation. In other words, excluding the small share of fallow parcels sitting very adjacent to irrigation would shift the reference baseline only modestly, not overturning our results. However, in application, readers should keep this proximity and spillover effect in mind: parcel-level applied-ET estimates for individual fallow fields near irrigated neighbors may still carry a meaningful upward bias.

Figure 6: **Robustness of the monthly applied-ET trend to the choice of satellite ET product, Boser (blue), Jacobs (red), and statistical subtraction (green), for (a) California and (b) Utah, April-October.** Points show the field-month mean applied ET for each method and error bars span the interquartile range; panels facet by crop.

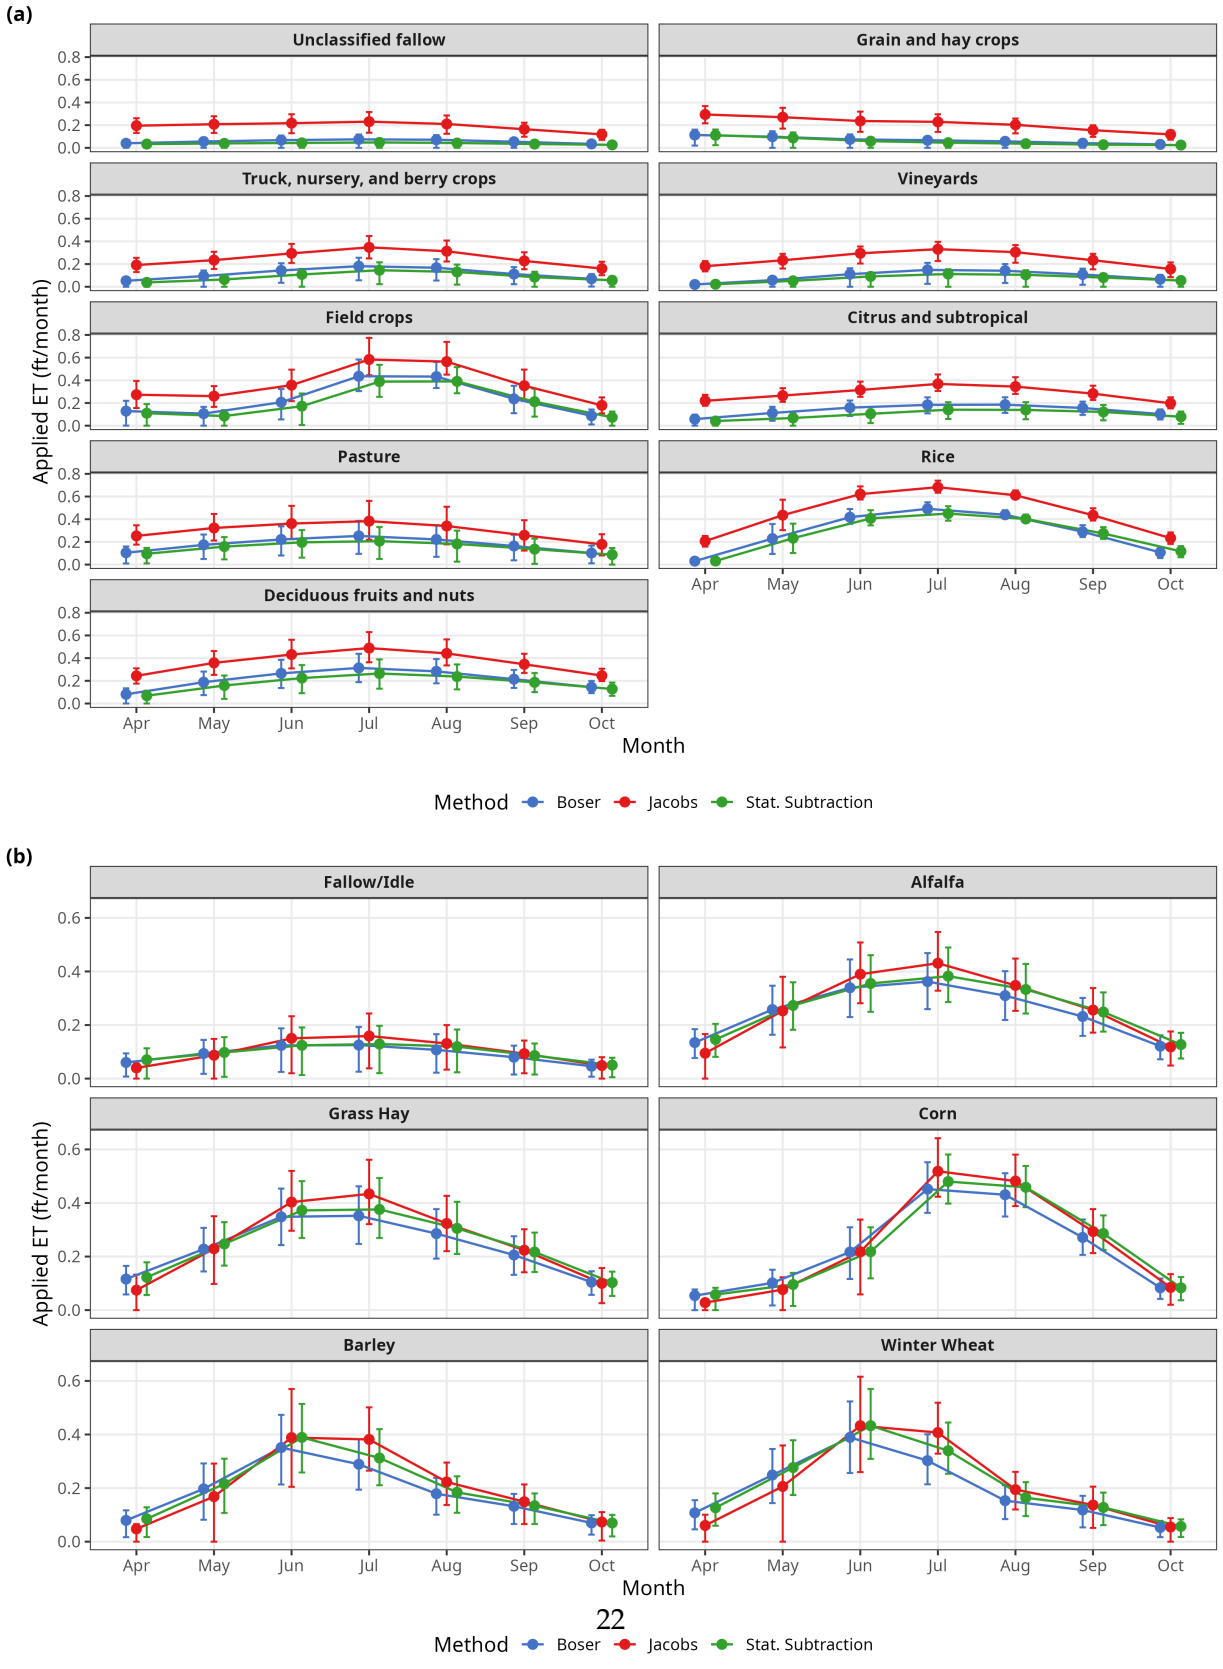
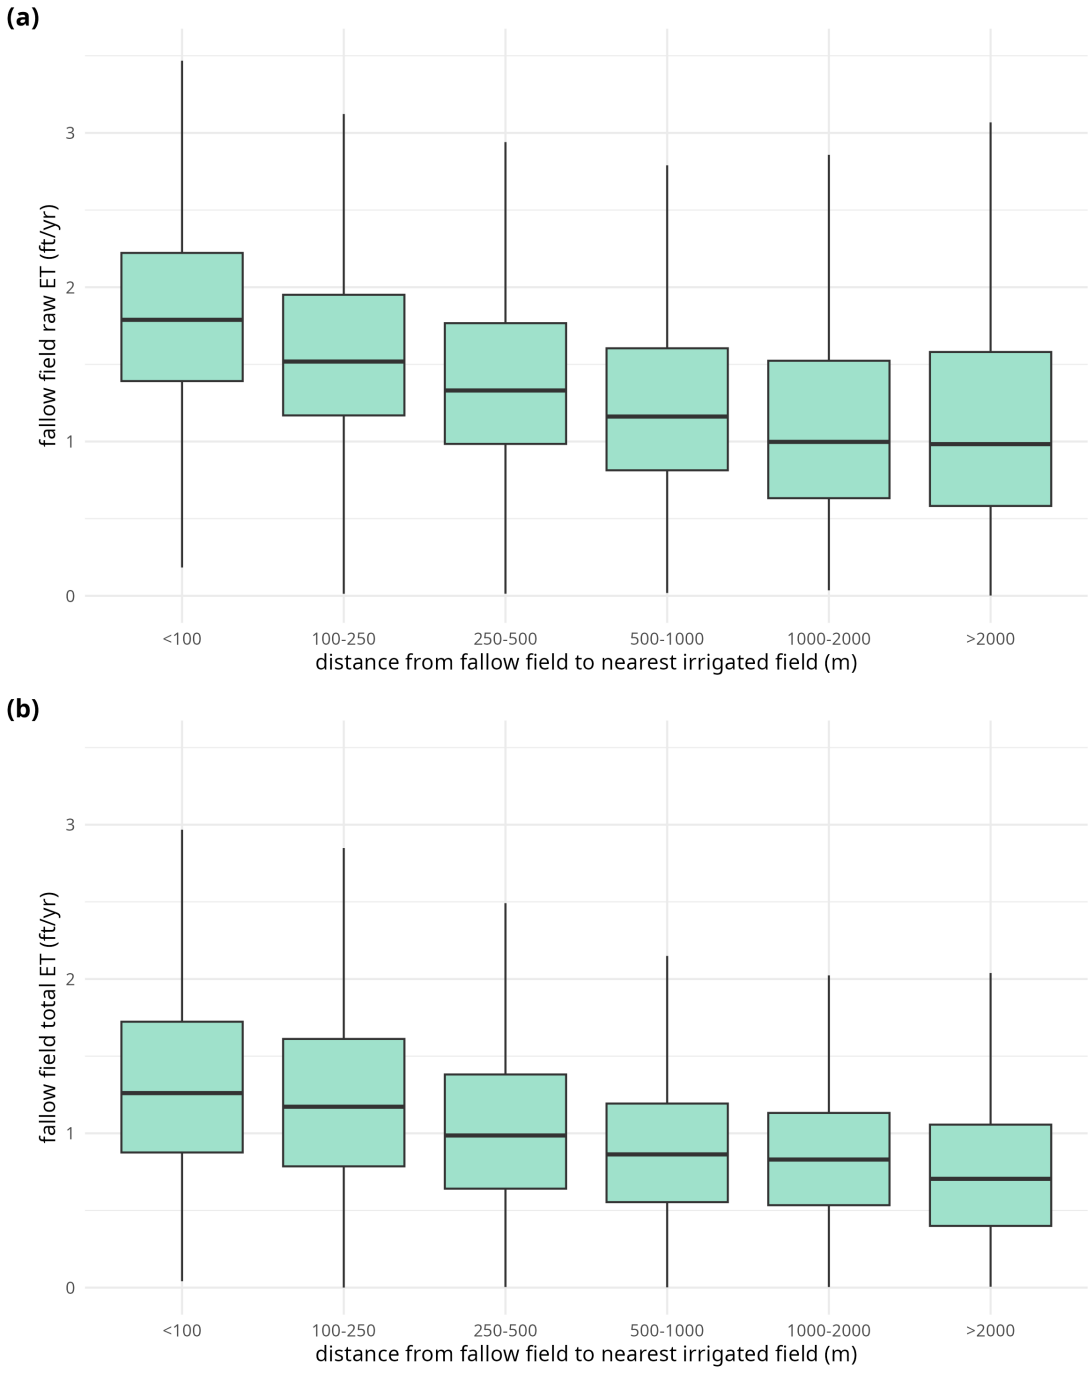


Figure 7: **Counterfactual-selection check for the fallow/idle reference parcels used to benchmark naturally occurring ET, for (a) California and (b) Utah.** Each fallow field-year is matched to its 8 nearest irrigated fields in the same year, and boxplots show the fallow parcel's annual ET (actual ET) binned by distance to the nearest irrigated neighbor. A declining relationship between fallow ET and distance would indicate that fallow parcels near irrigated land receive incidental water (runoff, lateral seepage, or misclassification) rather than representing a genuinely non-irrigated baseline, which would bias the applied-ET comparisons in Boser and Statistical Subtraction Methods toward understating true crop water use. Both panels show median fallow ET falling as distance to the nearest irrigated field increases, with the decline leveling off beyond roughly 1 km.



6 Conclusion

Satellite-based estimates of evapotranspiration are increasingly used to measure agricultural water use, yet relatively little is known about how these estimates compare with the crop coefficient methods that have traditionally been used to measure crop water requirements. This paper provides a systematic comparison of these approaches. Using field-level data from California and Utah, we compare Crop ET with satellite-based estimates of applied ET.

Our first contribution is to document where crop coefficient and satellite approaches agree and where they diverge. The differences vary across crops and settings. In Utah, for example, the approaches produce similar annual estimates for alfalfa but substantially different estimates for several other major crops. Examining monthly water use shows that these discrepancies often reflect differences in the assumed timing of water use rather than differences in levels. Crop ET tends to impose relatively concentrated growing seasons, while satellite ET frequently indicates flatter profiles and water use extending later into the season. We develop a novel approach for expressing these differences as empirical crop coefficients, allowing the seasonal water-use patterns observed by satellite to be compared directly with the engineering (K_c) curves.

Our second contribution is to evaluate how sensitive satellite-based estimates of applied ET are to the method used to separate irrigation-driven ET from naturally occurring ET. We compare the structured Jacobs water-balance approach with the computationally intensive machine-learning approach developed by Boser et al. and with a parsimonious Statistical Subtraction method based on spatial and temporal fixed effects. Despite their very different structures and information requirements, the three approaches generate surprisingly similar distributions and seasonal patterns of applied ET in Utah. This suggests that uncertainty about natural ET is not the primary source of the differences we observe between satellite and crop coefficient estimates. It also shows that relatively simple statistical methods can approximate much more computationally intensive approaches for recovering applied ET from satellite observations.

These findings have implications for the growing use of remote sensing in water accounting, program evaluation, and economic research. What approach is used matters for estimated levels and the timing of water use. Combining the structure and agronomic interpretation of Crop ET with the field-level information contained in satellite observations offers a promising path toward more accurate and empirically grounded measures of agricultural water use.

References

Richard Allen, L. Pereira, Dirk Raes, and M. Smith. Fao irrigation and drainage paper no. 56. *Rome: Food and Agriculture Organization of the United Nations*, 56:26–40, 01 1998.

Burdette Barker, Matt Yost, and Sena Bildim. Quantifying depletion differences from irrigation practice changes in utah, 2025. URL https://waterrights.utah.gov/Library/Misc_Reports/USU_20250513_QuantifyingDepletionDifferencesfromIrrigation.pdf.

Anna Boser, Kelly Caylor, Ashley Larsen, Madeleine Pascolini-Campbell, John T. Reager, and Tamma Carleton. Field-scale crop water consumption estimates reveal potential water savings in california agriculture. 15(1):2366, 2024. ISSN 2041-1723. doi: 10.1038/s41467-024-46031-2. URL https://www.nature.com/articles/s41467-024-46031-2.

Robert W. Hill, C.E. Brockway, R.D. Burman, L.N. Allen, and C.W. Robison. Duty of water under the bear river compact: field verification of empirical methods for estimating depletion, 1989.

Robert W. Hill, Burdette Barker, and Clayton S. Lewis. Crop and wetland consumptive use and open water surface evaporation for utah, 2011.

Rania Lachhab and Ellen M. Bruno. The value of water storage: Evidence from california's kern water bank. Working paper, July 2026.

Christopher Pearson, Blake Minor, Charles Morton, John Volk, Christian Dunkerly, Eric Jensen, Peter Revelle, Ayse Kilic, Richard Allen, and Justin L Huntington. Historical evapotranspiration and consumptive use of irrigated areas of the upper colorado river basin, 2024. URL https://rgdoi.net/10.13140/RG.2.2.31069.63206.

Isabelle Picciotto. A clean SWEEP? estimating the impacts of irrigation efficiency subsidy programs on water and energy use. Working paper, 2025. URL https://spaces-cdn.owlstown.com/blobs/sgny2gcxlkt85da9gyzxncn15btt.

Safoura Safari, Soroush Sharghi, Reza Kerachian, and Hamideh Noory. A market-based mechanism for long-term groundwater management using remotely sensed data. 332:117409. ISSN 03014797. doi: 10.1016/j.jenvman.2023.117409. URL https://linkinghub.elsevier.com/retrieve/pii/S0301479723001974.

USDA. *Part 623 National Engineering Handbook*, chapter Chapter 2: Irrigation Water Requirements. United States Department of Agriculture, 1970.

John M. Volk, Justin L. Huntington, Forrest S. Melton, Richard M. Allen, Martha Anderson, Joshua B. Fisher, Ayse Kilic, Anderson Ruhoff, Gabriel B. Senay, Blake Minor, Charles Morton, Thomas Ott, Lee Johnson, Bruno Comini de Andrade, Will Carrara, Conor Doherty, Christian Dunkerly, MacKenzie Friedrichs, Alberto Guzman, Christopher Hain, Gregory Halverson, Yanghui Kang, Kyle Knipper, Leonardo Laipelt, Samuel Ortega-Salazar, Christopher Pearson, Gabriel E. Parrish, Adam J. Purdy, Peter M. ReVelle, Tianxin Wang, and Yun Yang. Assessing the accuracy of openet satellite-based evapotranspiration data to support water resource and land management applications. *Nature Water*, 2:193–205, 2024. doi: 10.1038/s44221-023-00181-7.

Katherine Wright, Andrew Ayres, and Bryan J. Leonard. Mind the gap: Improving water markets with field-level remote sensing. Working paper, 2026. URL https://bryanjleonard.com/files/working-papers/WAL_PVID_2026.pdf.