Supplemental information for:

Shaping land use change and ecosystem restoration in a water-stressed agricultural landscape to achieve multiple benefits

Benjamin P. Bryant, T. Rodd Kelsey, Adrian L. Vogl, Stacie A. Wolny, Duncan J. MacEwan, Paul C. Selmants, Tanushree Biswas, and H. Scott Butterfield.

Correspondence: bpbryant@stanford.edu

This PDF file includes:

Tables and Figures labeled “SM” appear in the methods section while “SR” appear in the results section.

Methods: Supplementary details

Developing production and retirement statistics for futures with and without SGMA

We applied the most current version (as of Summer 2018) of the Statewide Agricultural Production (SWAP) model[1] to evaluate permanent and temporary land retirement under alternative SGMA scenarios. SWAP is a regional agricultural production and economic optimization model that simulates the markets for California crops. SWAP is linked to regional ground and surface water data and used to simulate the response of the agricultural sector to changes in groundwater availability and cost caused by implementation of SGMA. While the model structure and calibration procedure are both well developed (1, 2), it was necessary to specify alternative SGMA implementation scenarios, changes in water availability under climate change, and to conduct some additional post-processing to infer permanent and temporary land retirement and the overall impact of SGMA.

[1] http://swapmodel.com/detail

SGMA scenarios to represent in SWAP

We developed three scenarios in the SWAP model representing baseline future conditions without SGMA, future conditions under SGMA with significant investments in additional water capture and storage, and future conditions under SGMA with minimal additional water supply investment (Table SM1). Within the SWAP products, we refer to these scenarios as "Future without SGMA," (or "No-SGMA" for short), "Low Impact SGMA," and "High Impact SGMA," respectively. After feedback from stakeholders, we decided to focus on the "High Impact" scenario as most salient of the two with-SGMA scenarios, as well as to reduce the combinatorial expansion of scenarios considered when presenting the main results. The "High Impact" scenario is the "BAU" or "BAU SGMA" scenario referred to in the main manuscript.[2] (Even the high impact scenario shows more modest impacts to ag than existing studies.)

The first three rows of Table SM1 are constant across scenarios. Historical hydrology refers to the sequence of years used to represent variability in water supply by region. Climate conditions refer to changes in future water supply deliveries under climate change. A constant 2030 climate is applied to all scenarios over the same historical hydrology. SGMA is implemented as region-specific constraints that limit groundwater pumping to be no greater than the sustainable yield in each subbasin. The sustainable yield target is implemented by limiting total regional groundwater pumping based on an approximate percent difference in pumping above or below the sustainable yield, by hydrologic year type, such that the average annual safe yield is met over the representative set of hydrologic conditions. In addition, regional water supplies are augmented based on development of additional storage and investment in new, local infrastructure. The "Water Supply and Use" row summarizes these factors for each scenario.

Table SM1 -- Key scenario elements for the SWAP modeling.

	No SGMA **(Base at 2040)**	**High-impact (BAU) SGMA (at 2040)**	**Low-impact SGMA (at 2040)**
Climate conditions	2030	2030	2030
Historical hydrology[3]	1974 – 1994	1974 – 1994	1974 – 1994

[2] Certain intermediate products retain the original scenario label of "BASE" for "No-SGMA" -- this is not the same as "BAU" in the main manuscript, where BAU refers to what we consider the most likely, in contrast to a habitat-aware strategic coordination approach.

[3] Note that actual water supply estimates from the historical period are used for calibration, but for the future condition scenarios, modeled results for water years 1974-94 incorporate 2030 climate conditions as developed for the Technical Reference for the California Water Commission's Water Storage Investment Program.

Market conditions	Shifting consumer demand; endogenous prices; dairy market contraction	Shifting consumer demand; endogenous prices; dairy market contraction	Shifting consumer demand; endogenous prices; dairy market contraction
SGMA	No SGMA	Phased SGMA implementation to 2040; steady-state implementation at 2040	Phased SGMA implementation to 2040; steady-state implementation at 2040
Water Supply and Use	No additional storage	Additional water available for recharge based on existing infrastructure	Sites reservoir + CA “Water Fix + investment in additional recharge. Additional flexibility in farm practices to adjust irrigation.

SWAP primary and derived outputs

Average annual cultivation statistics

For each scenario and region, SWAP generates annual cropping statistics for eighteen different crops over a simulated 21 year period of representative hydrology (with water availability -- but not agronomic response -- adjusted for climate change). The model provides outputs that include region-specific estimates by year, which were then summarized into annual averages including gross revenue, net revenue, applied water, and consumptive water, for each of 18 crop categories, for each region, and for each of the three scenarios mentioned above. Our modeling of land use change aggregates crops to annuals and perennials, but when doing so it tracks the associated statistics based on area-weighted contributions of the finer-scale classification within each scenario and region. For example, the per-area returns to perennials in Region 1 versus Region 9 will reflect the fact that there is a different portion of almonds versus grapevines in the two regions.

Estimating permanent retirement by SWAP region

A decrease in average annual cultivation may arise from a mix of reduced cropping intensity on cultivated land and from permanent retirement of land currently in agriculture. This distinction is of interest because the two land types are associated with different opportunity costs when assessing whether to bring parcels into alternate land use (viz. restoration). In order to estimate a partition between these two sources of reduction, we calculated a change in region-specific *maximum cultivated area* within each of the SWAP scenarios, based on an average of the three highest-cultivation years in each scenario (typically, but not necessarily, corresponding to the wettest years). The difference in these maximums across the with-SGMA and no-SGMA futures represents an estimate of permanent retirement attributable to SGMA, while the difference between average annual cultivation and the physical footprint of ag represents an estimate of average temporary fallowing, with some nuances discussed further below.[4] Because there may be substitution between crops within the 21 year period analyzed and permanent retirement is not specific to a type of crop, these results are only provided at the region level for the two SGMA scenarios, relative to the no-SGMA scenario.

Consideration of uncertainty in SWAP

While we explore many scenarios and parametric uncertainties in the broad workflow of this paper (described below), we acknowledge these are essentially all conditional on the regional-level SGMA impacts as modeled by the difference between two focal SWAP scenarios. It would of course be desirable to undertake additional exploration of the SWAP assumptions and inputs, though we do not believe they would fundamentally change the results in surprising ways. First, our preliminary explorations of the resulting conservation portfolios under "low-impact SGMA" were relatively similar. Second, as mentioned earlier, the "High Impact/BAU" scenario is regarded by some in the SJV space as a low level of fallowing compared to some other estimates. To the extent this is correct and overall fallowing will be higher, this lowers the opportunity cost associated with restoration and essentially gives more scope to prioritize based on conservation. Furthermore, from a practical standpoint, improved fallowing estimates are likely to emerge in the coming few years as GSAs establish water budgets and sustainability plans, as well as positions on water trading that will fairly dramatically affect required fallowing levels. Nevertheless, we are currently beginning follow-on work to conduct a more in-depth sensitivity analysis of the feedbacks between the land use change modeling steps and the resulting conservation portfolios.

Modeling land use change to translate regional SWAP model output to the pixel level

SWAP outputs are provided for 10 regions within the SJV and so provide some coarse indication of the spatial variability of SGMA impacts. However, because habitat is fundamentally spatial with variation on a much finer scale, our workflow requires translating the available SWAP output into plausible patterns of retirement and cropping intensity on a scale that aligns with spatial variation in habitat. For this analysis, we use a 1080m pixel, though as described below we track fractional land use below the pixel level. We create these landscapes by implementing a rule-based algorithm that uses spatially varying agricultural suitability layers to modify the existing landscape to match region-level statistics from SWAP. Given numerous uncertainties governing land use dynamics, we develop a suite of *plausible* pixel-level spatial patterns of cropping in a future landscape, rather than focus on estimating a most-likely scenario. We first describe the rules represented, and then the spatially varying drivers.

Rule-based approach to generate BAU landscapes

Our rule based approach is implemented to capture the following qualitative dynamics:

---

[4] Terminology in SWAP products refers to the difference in average annual cultivation between with-SGMA and without-SGMA futures as "temporary fallowing" but we apply a separate definition more focused on physical land use, explained in the next section.

- Switching crop categories is costly. Therefore, all else equal, area dedicated to a crop category (ie, annual or perennial) that remains in production should be assigned first to places where it was already being grown in our base year landcover map.
- Suitability by definition means that, within a crop category whose area is being *decreased*, less suitable land will go out of production before higher suitable land.
- Suitability also means that, within a crop category seeing an *increase* in planted area, higher suitability land will be brought in first.
- Perennials have priority over annuals due to their higher returns.[5]

These rules together define a fairly straightforward mechanistic algorithm, expressed here at a high level, and applied within each SWAP region:

1. **Land cover stack:** Starting with a base year land cover map, reclassify all land in the base year to one of four categories: perennial agriculture, intermittently cultivated annual agriculture, natural, or excluded. Intermittently cultivated annual land includes land identified as idled, and is ultimately associated with a region-specific cropping intensity, which could be greater or less than one.[6] To ensure greater accuracy and avoid pixelization error, the area of each land class within each pixel is tracked (based on the finer scale 30m land cover), rather than assuming an entire pixel is of one particular class. The limits of the present-day agricultural footprint are taken as the limits of no-SGMA future, corresponding to an assumption of negligible expansion of footprint between now and full SGMA implementation, and also negligible permanent retirement in a no-SGMA future.
2. **Permanent retirement:** For a with-SGMA scenario, the SWAP analysis provides an estimated area to be permanently retired. To assign this to specific pixels on the landscape, agricultural land on pixels with the lowest values for the agricultural suitability index is decreased to zero on incremental pixels, until the region-specific permanent retirement value has been met. For most pixels being retired, this will involve retiring all annual and perennial cropland within the pixel. Except in cases of very rare numerical coincidence, each region will also have one "remainder" pixel that gets partially retired. In these cases, annual area is reduced first, and then if additional retirement is still necessary within the pixel, perennial area is further reduced.
3. **Perennial expansion or contraction:** The area assigned to perennials is expanded or contracted relative to the base year, so that the footprint in the scenario map matches the average annual outputs from SWAP for perennials. (This effectively assumes that the footprint of perennials is in a 1:1 ratio with the average annual values -- i.e., that planted area is in a steady

[5] There are two different factors involved in this assumption: One is that spatially varying factors affecting suitability are the same for annuals as for perennials. The other is that, even if they are very similar in their definitions of suitability, perennials will take priority -- which seems generally plausible due to their higher returns, though there are anecdotes where that is not the case, due to value chain considerations and local processing infrastructure.
[6] Our 30m base year landcover to which the above reclassifications are applied is included in the data upload. It is an augmented version of LandIQ/DWR (https://databasin.org/datasets/6cc5b24e401043a899a6db6eef5c86db), augmented with sources such as CalFire's FVeg (https://map.dfg.ca.gov/metadata/ds1327.html), OpenStreetMap for roads and rail, and NASS Cropland Data Layers (CDL) from 2014, 2016, 2018 to handle ties regarding some conflicting or incomplete land uses. We use LandIQ as the base layer because it is generally regarded as of greater thematic accuracy than CDL. Landcover crosswalks are listed in the included file "lulc crosswalk master.xlsx"

state after SGMA implementation, and trees are not being taken into or out of production each year, as that is agronomically not possible.)

a. When perennials expand, they are assumed to expand on the highest quality ag land that is not under already under perennials and was not permanently retired in Step 2 -- in effect displacing land that would be under intermittent cultivation of annuals (but note that this a statement about *long-run land use*, not planted area -- i.e., expansion or contraction of perennials does not imply change in average annual production of annuals, just the extent of the footprint on which it occurs, as described in Step 4).
b. When perennials contract, they are assumed to contract from the lowest valued pixels on which they are grown, in essence “freeing up” that land for intermittent cultivation.

4. **Annuals are assumed to be intermittently cultivated** over the remaining (unretired) agricultural land within the region, with their average intensity of cultivation determined based on the SWAP average annual production within the region, divided by the footprint available for cultivation of annuals.
5. To create maps that describe the (revenue or water use[7]) **impact of retiring agricultural land** on a given parcel, we apply the following assumptions:
    a. Land under perennials receives the per-area region-specific value from SWAP (weighted by area of the sub-categories of perennials). E.g., if SWAP output indicates that, on an area-weighted basis, perennials have a gross revenue of 2000 $/ha (per year), and the pixel in question contains 50 ha of land under perennials, the revenue layer is assigned a revenue value of 50 x 2000 = $100,000 for that pixel.
    b. Land under intermittent annuals receives the per-area region-specific value from SWAP (also weighted by area of the sub-categories of annuals), and scaled by the cropping intensity of annuals within the region. For example, if the region in question contains 10,000 ha of land dedicated to annual crops, and the average annual production is 6,000 ha, with a per-hectare return to planted area as $1000/ha, this implies that an average hectare of land under intermittent annuals would have a gross return of $600/ha. If a pixel contained 50 ha of land under intermittent annuals, it would be assigned a revenue value of 600 x 50 = $30,000. This corresponds to the assumption that all land in the footprint of intermittent cultivation is equally likely to be cultivated and that the cost of retiring that land us uniform within a region.
    c. Finally, since area within pixels are tracked separately for annual and perennial, the two are added together within the pixel.

With the exception of the initial land cover reclassification, the code to implement this algorithm is in the file RB_downscaling.R

Estimating temporary fallowing

Temporary fallowing as reported in the main manuscript is interpreted as the average annual area of land within the (unretired) agricultural footprint that is *not* planted. SWAP is based on aggregate production statistics and does not take physical land use as an explicit input or output. Therefore, to estimate temporary fallowing as we define it, we apply conversion factors to translate *average annual production statistics* to *physical net area planted* in an average year. Corn is the dominant crop that sees multiple plantings per year, and so we apply the following conversion factors to the corn acreages within the

[7] Water use is not currently considered in the analysis presented in the paper but is mentioned because the code implements this as well, to preserve possibilities for future analysis.

SWAP outputs. These factors were developed from earlier DWR and county data estimates of double-cropped silage acreage.

Table SM2 -- Net cropping area correction factors.

Correction factors applied to the corn category within SWAP outputs, to translate from average annual area planted to physical footprint.

Region	Silage correction factor to translate to net area planted
1	60%
2	60%
3	63%
4	60%
5	63%
6	42%
7	43%
8	43%
9	54%
10	63%

Average annual temporary fallowing is then calculated as the gap between the total of (permanent retirement and average annual net area planted) and the area of the present-day agricultural footprint. Or, equivalently, the gap between future agricultural footprint, and future average annual net area planted. Mismatches between SWAP modeling assumptions and land use assumptions result in a small negative temporary fallowing in Tule (Region 8), which we report as zero for logical consistency (negative temporary fallowing does not have a meaning). This may be due to three sources of error: 1) Overestimate of total production statistics by SWAP-RTS, 2) Overestimate of net-area conversion factors for corn or other crops, 3) Underestimate of physical footprint in agricultural land.

Additional considerations and assumptions

1. Our approach does not take into account expansion of urban lands or other land use change (solar, etc), which are assumed static. Future work could incorporate expansion of other land uses as constraint layers on where restoration can be undertaken. This approach would not require fundamental changes to the workflow, but would still not incorporate feedbacks related to spatial variation in land scarcity and competing demands -- that would require fundamentally different modeling approaches to land use and opportunity cost.
2. The connection between SWAP results and land use assumes the boundary of present-day agriculture is the boundary of agriculture in a no-SGMA future. While some areas may see small amounts of expansion, we expect that even in the absence of SGMA the variability of surface water availability under climate change, increased pumping costs, and environmental regulations would limit significant expansion relative to the present day. This assumption is generally most likely to be violated in the northern end of our study region, where we find minimal high quality

habitat for the target species -- so while it may affect land use projections, it is unlikely to significantly affect restoration targeting.

3. We assume that existing natural land will remain natural, regardless of whether it is currently protected. While we account for existing protected lands in our analysis, this exercise is not explicitly about prioritizing threatened natural land, it is about identifying priority land currently in agriculture that can serve habitat goals at minimal cost to the ag community -- but within a broader landscape of existing natural and protected lands. The locations of the ag lands that are high priority for restoration may incidentally suggest unprotected natural lands that merit protection to ensure improved contiguity or support migration corridors, but formal consideration of those factors is beyond the scope of this analysis.
4. Land under conservation easement is treated as both protected from development and, if it is agricultural land, is modeled as protected from retirement. As background, LULC maps indicate that existing protected areas (ie, those in some type of conservation) cover some land that has non-negligible amounts of ag in the present day. In addition to being treated as protected from development (the intuitive conception of "protected"), *existing ag on existing protected land* is treated as *protected from ag retirement* in the transition from the current landscape to a BAU future.

Creating multiple plausible ag suitability layers

The above algorithm relies on a spatially defined agricultural suitability index that is generally reflective of spatial variation in anticipated returns to agriculture, in an ordinal sense: That is, the lowest values represent the "worst" land that is therefore most likely to be permanently retired, and lower and higher values generally correspond to places where growers would prefer to avoid or expand, respectively. The cardinal values of the agricultural suitability index are *not* assumed to have meaning: E.g., an agricultural suitability index of .8 is not assumed to translate to 33% higher yields or profit than land with an agricultural suitability index of .6, or to 33% greater likelihood of being cultivated, etc. More sophisticated and data-intensive strategies for developing a suitability index could be utilized to speak to those interpretations, and in turn inform the cost surface as well, but that is beyond the scope of this study.

To generate realizations of an ag suitability index, we consider linear combinations of spatially explicit factors like soil quality, salinization status, and security of access to water supply, based on the assumption that they are predictive of higher and lower ag suitability. By combining a broad set of factors that are reasoned to be relevant based on contextual knowledge (elaborated in the following section) and doing so using a range of weights, we generate a range of plausible agricultural suitability patterns.

Specifically, we generate multiple agricultural suitability layers by combining different weight combinations on a set of spatially explicit factors (below), with the final ag suitability index value (AGSI) for a given layer taking the form of a normalized weighted index:

$$AGSI = \beta_1 F_1 + \beta_2 F_2 + ... + \beta_K F_K; \quad \sum_k \beta_k = 1$$

Where each F_k is a potential spatially explicit driving factor (hereafter "ag suitability factor") such as those referred to above, but which have been normalized so that when a factor takes a value of 1 in a pixel this indicates it is the best possible value for agriculture, and when it takes a zero it is the worst possible value. Weights themselves also vary between zero and 1, so that setting a weight to 1 on any particular factor and zeros on the others equates to answering the question "what if that particular factor was the sole predictor of ag suitability?"

After exploration of a larger number of scenarios, we focused on four spanning scenarios and one evenly weighted scenario as shown in Table SM3.

Table SM3 -- Weights for agricultural suitability factors.

Spatially explicit suitability factors and weights determining different ag suitability layers (columns), with cells indicating the weights places on each factor, and the corresponding short-hand name associated with each weights combination.

Driver layer → ------------------- LUC Scenario Name ↓	Land Assets	Land Impairment	Groundwater	Surface Water
Even	1/4	1/4	1/4	1/4
Land Assets	3/4	1/12	1/12	1/12
Land Impairment	1/12	3/4	1/12	1/12
Groundwater	1/12	1/12	3/4	1/12
Surface Water	1/12	1/12	1/12	3/4

Sources and development of agricultural suitability factors

We draw our ag suitability factors from two sources: Some land quality and impairment indices provided by another study of the future agriculture in the SJV (3), and some custom-developed indicators of surface and groundwater security built by integrating publicly available water data.

Land Assets and Land Impairment

The American Farmland Trust and the Conservation Biology Institute (CBI) recently undertook a similar but distinct exercise as ours, to identify spatially varying drivers of high quality or high importance farmland for protection (3). While having a somewhat different orientation, this study ultimately assembled and combined various data layers designed to indicate spatial variation in factors relevant to ag suitability. It includes a Land Assets layer that relies on the California Storie Index, Farmland Mapping and Monitoring Program Rank, Aquifer Recharge Potential, and presence of Citrus Microclimates. It also includes a Land Impairment layer that integrates spatial data on saline soils, sodic soils, shallow groundwater tables (having potential for waterlogging), and land fallowed during the 2012-2017 drought. In each case the constituent layers were combined using CBI's Environmental Evaluation Modeling System,[8] which integrates the layers using a tree-based fuzzy logic approach. For our analysis, we used the two high level land asset and land impairment layers for the land-quality related drivers.

[8] https://consbio.org/products/tools/environmental-evaluation-modeling-system-eems

Custom factors for surface and groundwater risk

We developed spatial indicators intended to capture where surface and groundwater supplies were likely to be more or less reliable and more or less constrained in a water-scarce future. In doing so, we considered historical predictors and factors that, when combined with the logic of SGMA implementation, may be predictive of greater or lower likelihood of cultivation or retirement. A key challenge in this regard is being aware of how historically valid predictors of ag suitability may change under future policy and climatic conditions. For example, under open access to groundwater, historical groundwater overdraft may be highly predictive of high value crops, but will actually be an indicator of lower water security in a future for those lying within Groundwater Sustainability Agencies (GSAs) where that groundwater overdraft is halted -- suggesting that statistical inference based on historical features may not be suitable.

Surface water security index

Our surface water security index is based on a combination of three factors:

- Relative security of surface-water rights, as a ratio of **estimated present-day water demand** to a measure of **secure water rights**;
- **Historical fallowing** patterns as a proxy for insecure access to water during droughts.

The index itself is calculated by combining the fallowing dataset and the ratio of applied to water to the estimate of secure water rights.
$V_s = w_1(S/S_{max}) \ + \ w_2((NF)/(NF_{max}))$
with
$S = log\ ((R\ +1)/\ (A\ +\ 1))$

where:
$S =$ ratio of applied water to total pre-1914 water rights that supply it
$S_{max} =$ maximum value of S within each region
$V_s =$ surface water vulnerability index, linearly scaled to range [0, 1]
$A =$ total current water demand (applied water) in the GSA (acre-feet/yr)
$R =$ total claimed pre-1914 water rights for diversion points reported in eWRIMS for the GSA (acre-feet/yr)
$NF =$ number of years fallowed [0, 5]
w indicates relative weights on the applied water versus fallowing factors. For these study they are taken as relative weights of ⅔ and ⅓.

The constituent layers are described below:

Surface water available as senior water rights by GSA (R)
Senior water rights are an indication of the ability to secure water by individuals or irrigation district when surface water supplies are low. We rely on data from the electronic Water Rights Management System (eWRIMS) maintained by California's State Water Resources Control Board to determine access to senior water rights at the GSA level.[9] eWRIMS is a public database providing basic information on water rights, including locations of points of diversion (POD) associated with each water rights application, owner name, status, date, pre-1914 and riparian designation, diversion amount, and beneficial use (e.g.,

[9] https://ciwqs.waterboards.ca.gov/ciwqs/ewrims/EWMenuPublic.jsp

hydropower, agriculture, domestic, industrial, recreation, and environmental). The database provides the diversion *amount* only for appropriative (post-1914) water rights, and provides only a diversion *rate* for pre-1914 and riparian rights. Importantly, none of the rights in this database directly account for available water supply in a given year, and surface water is generally recognized as being overallocated (4); nevertheless their variation in space should still provide some relative indication of greater and lesser water security, and this relative status is the only information we utilize in our algorithm.

We use the annual water diversion rate assigned to pre-1914 surface water rights claimed for points of diversion located within each GSA as the basis for access to senior water rights, which informs the normalization. For example, if there are three GSAs within a subbasin, the GSA with the lowest surface water rights would receive a zero, the GSA with the highest senior water rights would receive a one, and the middle GSA would receive a value between zero and one based on where it's pre-1914 diversion rights sit relative to the two extremes.

Although water usage reports claiming pre-1914 rights that have been filed with the State Water Board are admittedly incomplete (4), these are currently the best data available on the availability of the most senior water rights in the system. While riparian rights are also senior, they are legally distinct from pre-1914 rights in that 1) they are not allowed to be transferred to other lands beyond the riparian property, and 2) they are only realized when there is water available to divert. In contrast with pre-1914 water rights, therefore, we do not consider riparian rights as secure a source of water since they are more vulnerable to year-to-year fluctuations in river flow and diversions upstream. There is also regional variability in reliance on imported water and the security of imported water supply from the Central Valley Project, but we were unable to collect "wall-to-wall" data on this. Therefore it is omitted from the current analysis, but would be helpful to include in future work.

Applied Water (A)
DWR provides agricultural water use information by crop type per county[10], and the average applied water by crop type and county was calculated over the years 1998-2005. These average values were mapped to the DWR/LandIQ land cover map (at 30m resolution). We then summed the pixel-level applied water values for all pixels within each GSA, resulting in a final applied water raster with units of acre-feet per year per GSA. Since 2014 was a drought year that was already demonstrating the effects of low surface water availability, we assigned idled land the weighted average applied value associated with annuals, rather than zero -- using zero would have lead to an "artificially" flat demand surface as if demand had already been curtailed.

Number of years fallowed (NF)
The number of years fallowed is derived from NASA Fallowed and Cropped land in the Central Valley: https://nex.nasa.gov/nex/resources/370/.
Working citation: Melton et al. 2017. Satellite Mapping of Fallowed Agricultural Lands during the California Drought.

Data was provided directly by Melton for 2011 to 2016, with the caveat that 2012 data was notably less reliable,[11] so our analysis excludes that year. The layer developed does not distinguish non-ag land from land that was never fallowed during that period. We therefore assign zeros to missing pixels that are

[10] https://water.ca.gov/Programs/Water-Use-And-Efficiency/Land-And-Water-Use/Agricultural-Land-And-Water-Use-Estimates
[11] Melton, personal communication 2017-08-29

identified as ag land from our our LULC, resulting in a raster where each pixel corresponding to ag land in our LULC is assigned a value between 0 and 5 representing the number of years that agricultural lands were not in production between 2011 and 2016, and excluding 2012.

Groundwater security index

We assume that restrictions on groundwater pumping under SGMA will have the greatest impact in places where the historical rates of groundwater decline are highest, and where there is greatest dependence on groundwater supplies to meet water demands. Therefore our groundwater security is based on just two factors:

- Historical rate of groundwater decline.
- Relative dependence on groundwater supply.

Calculating the index of ground water security

For each pixel, we calculate an index of overall groundwater vulnerability by scaling the rate of groundwater decline (D) by the dependence on groundwater supplies (P). The result is then normalized between 0 and 1 based on the minimum and maximum values of D taken over the SWAP regions:

$V_g = (-DP - D_{min})/(D_{max} - D_{min})$

where

$V_g =$ groundwater vulnerability index, linearly scaled to range [0, 1]

$D =$ rate of groundwater decline (ft yr^{-1})

$P =$ fraction of total water supply from groundwater [0, 1].

These is then converted to an index of groundwater security (S_g), by taking the additive inverse:

$S_g = 1 - V_g$

Detail on the two layers is below:

Rate of groundwater decline (D)

We characterize each GSA with the mean rate of decline in spring groundwater levels, based on CASGEM's Groundwater Information Center[12] data on well levels at 1174 points in our study area, with data collected from spring of 2002 through spring of 2017. All of the monitoring points have data spanning at least 10 years during this period, with some (212) having data that span 15 years. We calculated the rate of decline for each monitoring point, by taking the difference between the earliest and the latest available groundwater level data and dividing by the number of years in the record. We then calculated the mean rate of decline for each GSA and for each sub-basin, by averaging together the rates for all points within each GSA and each sub-basin. The mean rate of decline within each GSA is assigned to each pixel within that GSA; pixels within GSAs which have no data points within their boundary were assigned the mean rate of change for the sub-basin. This method is subject to some potential bias based on the location and density of wells, but we are unaware of a priori reasons such bias would exist, and consider it reasonable to assume that such bias will be offset by our consideration of alternative weights and, we later intend to explore alternative sources for groundwater risk. Spatial resolution of groundwater

[12] http://water.ca.gov/groundwater/gwinfo/

decline could be refined in future work by connecting to groundwater models and GSA-specific deficits that will become more clear as groundwater sustainability plans are finalized.

Dependence on groundwater supplies (P)
We used data on the percentage of total water supply provided by groundwater for each subbasin, provided by CASGEM's Basin Prioritization study (5). We assign this value, expressed as a fraction between 0 and 1, to each pixel within the subbasin.

Conversion of base files into normalized suitability layers is implemented with comments in the file ag_suit_preprocessing.R, including documentation of adjustments and imputations for missing data and border effects.

Additional considerations

Depending on GSA governance choices and infrastructure development over the period of SGMA implementation, surface water reliability and groundwater vulnerability may be relevant to the GSA or subbasin mostly in aggregate, or they may remain very independent concepts that drive spatial differentiation. For example, if GSA's have good surface water conveyance throughout the GSA and choose to facilitate intra-GSA trading and also credit "in-lieu" recharge to allow trading between surface water use and groundwater pumping, the overall risk within the GSA is more homogenous, reflecting average surface and groundwater conditions. If access to surface conveyance is very heterogeneous within a GSA, and the GSA members also choose to keep groundwater and surface water management fairly separate, then differences in the reliability of access to surface and groundwater within the GSA could drive spatial cropping patterns to a greater extent than is reflected in our analysis. In our case, the only water-risk variation we consider below the GSA level is pixel-level fallowing, which in reality is likely to reflect a combination of water access and land characteristics, and has some slight covariation with the Land Impairment layer, since fallowing factors into that layer as well.

Spatial Optimization

We utilize the minimum set objective framing[13] in *prioritizr*, using a zones framework to differentiate between restoring uncultivated land and active retirement and restoration. Existing natural land is taken as a "locked-in" form of "restore uncultivated".[14] The step by step specifications are detailed in thoroughly commented file "optim_batching.R". Here we review the high-level assumptions embedded in the optimization:

1. There is a threshold for the fraction of a pixel that is identified as natural land that must be met or exceeded for that pixel to be counted as pre-existing "natural" by the optimizer (and therefore worth trying to cluster next to through the application of a boundary penalty). That threshold by default is 75%, but is explored in the parametric sensitivity analysis.
2. To be eligible for restoration (whether active or of BAU-retired land), a pixel must have sufficient restorable land such that after restoration, it would be above the threshold for being considered natural. Ie, existing natural + [potential] restored must be bigger than whatever threshold one has set for natural. Eg, if the natural threshold was 75%, and a pixel was 40% natural, then a pixel with another 40% ag would be eligible for restoration (because after restoration, the pixel would

[13] https://prioritizr.net/reference/add_min_set_objective.html
[14] https://prioritizr.net/reference/add_locked_in_constraints.html

be 80% natural), while a pixel that was only 20% ag would not (because after restoration it would only be 60% natural, which is still below the 75% threshold).

3. However, if restoration action is taken in a pixel, the assumption is that all restorable land within the pixel is restored, not just up to the 75% threshold. This is a proxy for feasibility, transaction costs, and the like.
4. A pixel that exceeds the natural threshold to begin with is not eligible for additional restoration – this is also a proxy for engagement/transaction costs: While it might be worth it to move something from 76% natural to 95% natural all else equal, actors would rather focus on areas that are going to get larger gains. At a threshold of 75% natural, it is likely that there is already a good deal of contiguity of natural land within the pixel so it will probably be able to provide connectivity to the surrounding pixels, and is therefore sufficiently natural to form part of an augmented reserve.
5. To encourage clustering, the optimizer rewards reducing the boundary length around land that meets at least one of the following criteria:
    i. Is classified as existing natural
    ii. Is classified as existing protected, regardless of what fraction of the pixel is actually "natural" according to an LULC.
    iii. Is selected for active retirement and restoration
    iv. Is selected for restoration of BAU-retired land.
6. Restoration (whether of BAU-retired or active retired) is limited to land that is high quality habitat for at least one of the five target species. This is not the case for existing natural land or for existing protected land, which are assumed to be "locked in". This assumption essentially splits the difference between an approach where clustering is valued on all natural/restored land (even if it's not high quality habitat for a target species), and one that only values clustering on high quality habitat. Instead, it enforces the notion that "making clusters of natural land is valuable, but only if done so by augmenting with high quality habitat."
7. The cost of actively retiring and restoring ("retire and restore"/"RR") ag land is the gross-revenue associated with that land class in that region. Each region has a different value for the cost of retiring perennials [based on the SWAP cropping mix of perennials in that region], or for intermittently cultivated annuals [based on the SWAP cropping mix of annuals, as well as the total average annual cultivation relative to total area of available for annuals – ie, the cropping intensity].
8. The cost of restoring BAU-retired land ("Restore Uncultivated"/"RU") is specified relative to the costs of the lowest-return ag lands -- that is, we want it to be less than the cost of retiring any land that is under cultivation. At the same time, since taking that land out of the "option space" for ag has an impact to the SJV economy beyond just the boundary of the SWAP region where the pixel is, the spatial variation in the costs to the ag economy should perhaps not be quite as stark, and could be balanced by a measure of the basin-wide opportunity cost of land. We represent this dynamic as a weighted combination of the lowest return to cultivated land within each region, and the basin-wide average return to intermittently cultivated annuals – this weighted combination is then scaled down by some fraction (*costfrac*, in eq below) representing the opportunity cost to ag of taking that land out of the "option-space" even after it was permanently retired. That default fraction is 10%.

$$costfrac\ [basinwt\ *\ ave\ per\ ha\ returns\ to\ ICA\ in\ SJV\ +\ (1-basinwt)\ *\ min\ per\ ha\ returns\ in\ region]$$

For narrative simplicity, the main scenarios presented in the paper assume zero weight on the basin-wide averages, after confirming that the qualitative conclusions were robust to an even-weighting.

9. Costs are only used to determine relative costs in space, and relative to a boundary penalty – i.e., gross revenue is only a proxy for ag impact. Not only is gross revenue a coarse proxy for long-term opportunity cost, we also know that farmers, GSAs, and conservation actors might work to shift some cultivation in space rather than scale down total cultivation in direct proportion to land that is retired and restored (and indeed, the point of strategic engagement is to help bring about these shifts). Such shifting or changing of cropping intensity will have a cost, but not the same as actually retiring the land. For example, if there are 10 (equally sized) parcels, and on average six are cultivated in any one year, then the average return on one parcel (and the average assumed cost of retirement) is 6/10 x [the return on continuous cultivation of one parcel]. If the optimizer/conservation actors identify two of those parcels that are high priority for restoration, the optimizer will model the loss as 2 * [6/10 * return on continuous cultivation of on parcel of land]. In reality, the remaining eight parcels may be farmed more intensely as a result of restoration on the target pixels, which would significantly (but not 100%) offset the loss from taking those two parcels out of production. These subtleties are *not* represented in our optimization, but are given context by the land use statistics in our bar plots.
10. This analysis assumes that ag land can be fully restored to the same habitat quality level as natural land, and that the time required to bring about this equivalent habitat quality is irrelevant, so that restoring land gets full "credit" towards being high quality habitat regardless of the land's history in agriculture, or other aspects of land condition – the assumption is all of those factors are built into the species distribution model.
11. Relatedly, the cost of physical act of restoring, and to the extent necessary, managing, restored land is not taken into account. The focus is on cost as perceived by the ag community in terms of loss of production or option value.
12. We assume restoration itself has no impact on water use, compared to land that is simply retired – in reality it may require a initial irrigation or other inputs to establish restoration vegetation, and have a potentially different ET profile than abandoned land, but the assumption is that this is negligibly *different* from retired land.

Note to facilitate interpretation of code

In earlier project stages, we also explored alternate problem forms. In one alternate, the optimizer could also select additional lands for retirement to secure water for chronically under-delivered water-dependent wildlife refuges in the SJV. This formulation spatializes the water saved from retiring a given pixel in the same manner as revenue is spatialized for the primary problem structure, with the additional option that retired lands can either be restored (the main problem specification), or retired to redirect water and with end use of the land left unspecified. We also explored a specification that would allow reconfiguring the landscape with recommendations about where ag in high value habitat areas could be shifted at least cost. We do not consider either of these actionable or relevant problem structures at this point due to a combination of higher requirements for data and model fidelity, and nuanced issues of implementation feasibility. They are mentioned here because structures to facilitate that analysis are retained in the code, with values zeroed out where relevant to ensure a problem structure that is equivalent to the simpler structure described in the main manuscript.

Excess nitrogen assessment

We approach the co-benefit of reducing excess nitrogen in the landscape by calculating two metrics: (1) the change in nitrogen inputs to the system due to altered land use (tonnes/region; with the dominant effect stemming from retirement), and (2) mean historical NO_3-N concentrations in domestic and public water supply wells, as a contextual proxy for potential impacts on populations in the affected area.

Excess nitrogen application

First, we estimate the excess nitrogen applied across No-SGMA, BAU-SGMA, and optimized scenarios. We calculate excess nitrogen based on fertilizer application rates from published studies (6–10) and alternate sources such as COMET-Planner documentation[15] and USDA's Agricultural Chemical Usage statistics for 2015 (Table SM4).[16] We also estimated nitrogen use efficiencies by crop type from (7, 10–18). We aggregated values to coarser categories but weighted by region-specific SWAP cultivated areas -- thereby reflecting regional differences in crop mixes. We assume that the difference between applied nitrogen and harvested nitrogen (calculated based on nitrogen use efficiencies) represents the "excess" nitrogen that holds potential to threaten human health via atmospheric or water-borne pathways. By differencing across scenarios, we can then also develop a metric for "excess nitrogen avoided." More detailed treatment of nitrogen fate and transport is complicated by the fact that there is a high degree of spatiotemporal heterogeneity in soil N cycling (10), the treatment of which is outside the scope of this study.

Nitrogen application rates were originally developed for the set of crop classes defined in the DWR/LandIQ land cover map, and later mapped to their equivalent SWAP crop category. In some cases, the mapping from DWR to SWAP was direct (as for cotton), in other cases, several DWR classes were combined into a single SWAP class (such as OTHDEC.) Where multiple DWR classes were combined, final SWAP values were calculated as an area-weighted average of the N application rates for all DWR classes included in the equivalent SWAP category.

Table SM4 -- Sources for nitrogen application rates by crop category.

SWAP Crop Category	Weighted DWR classes included	Nitrogen application source
ALFAL	Alfalfa; Miscellaneous grasses	Swan et al 2015
ALPIS	Almonds; Pistachios	Tomich et al 2016
CORN	Corn, Sorghum and Sudan	Tomich et al 2016
COTTN	Cotton	Tomich et al 2016
CUCUR	Melons, Squash and Cucumbers	Tomich et al 2016

[15] http://comet-planner.nrel.colostate.edu/COMET-Planner_Report_Final.pdf
[16] https://quickstats.nass.usda.gov/

DRYBN	Beans (Dry)	Tomich et al 2016
GRAIN	Wheat	Tomich et al 2016
ONGAR	Onions and Garlic	Tomich et al 2016
OTHDEC	Apples; Cherries; Peaches/Nectarines; Pears; Plums, Prunes and Apricots; Walnuts	USDA NASS, Tomich et al 2016
OTHFLD	Miscellaneous Field Crops; Miscellaneous Grain and Hay; Sunflowers	USDA COMET
OTHTRK	Bush Berries; Carrots; Cole Crops; Lettuce/Leafy Greens; Peppers; Strawberries	Tomich et al 2016, Perry et al., 2002
PASTR	Mixed Pasture	USDA COMET
POTATO	Potatoes and Sweet Potatoes	Tomich et al 2016
PRTOM	Tomatoes	Tomich et al 2016
RICE	Rice	Tomich et al 2016
SAFLR	Safflower	Munier et al., 2011
SUBTRP	Avocados; Citrus; Kiwis; Olives; Pomegranates	USDA NASS, Tomich et al 2016, Ayars and Phene, 2011
VINE	Grapes	Tomich et al 2016

Table SM5 -- Excess nitrogen applied by crop category.

The values are stored in the repository file “nitrogen_with_SWAP_crop_mapping_sheetonly.csv”.

SWAP Crop Category	Excess N applied (kg/ha)
ALFAL	2.3
ALPIS	151.6
CORN	158.2

COTTN	132.4
CUCUR	111.8
DRYBN	63.8
GRAIN	164.7
ONGAR	165.0
OTHDEC	102.1
OTHFLD	69.8
OTHTRK	180.0
PASTR	69.6
POTATO	144.0
PRTOM	143.7
RICE	95.2
SAFLR	69.8
SUBTRP	37.1
VINE	32.0

Vulnerability to excess nitrogen

Vulnerability to excess nitrogen is determined relative to groundwater supplies, with a focus on drinking water supplies. Vulnerability of drinking water supplies to nitrogen contamination is based on historical NO_3-N measurements in monitored domestic and public water supply wells. We obtained well locations and monitoring data from the California Water Board's Groundwater Information System (GAMA). The data are available from 2000 to 2018, but the wells monitored are not consistent from year to year, making trends for individual well locations or sub-basins challenging to reliably characterize. Therefore, for each subbasin, we took the mean of NO_3-N for all data points (representing all measured domestic and public water supply wells in the sub-basin) and across all years. This is of course an approximation that will obfuscate time trends and sub-regional variation, but we determined that the combination of data quality and complexity of fate and transport meant that a finer-scale analysis was not feasible. For context, the maximum contaminant level (MCL) for NO_3-N is 10 mg/L (California Code of Regulations 22 CCR §63341).[17]

The two subbasins corresponding to Region 10 of our SWAP analysis (Pleasant Valley and Kettleman Plain) did not have well records available through the GAMA system, so their results are not included in the overlays of excess nitrogen and vulnerability. These are contiguous subbasins that are outside the valley floor, and small in size and population. Additionally, where regions represented combinations of multiple subbasins, the region was assigned the simple arithmetic mean taken over the subbasins within the region.

When interpreting these results, it should be noted that even in areas where nitrate contamination is not yet a concern, nitrate levels in groundwater may continue to rise after N application is reduced, due to the fact that travel times from sources of nitrogen to water supply wells can vary from years to decades, and even to centuries for the deepest production wells (19). Reductions in excess nitrogen applied may not result in immediate reductions in observed nitrate levels in drinking water, and the social benefit of those reductions may not be directly proportional to the current severity of the problem. It is possible reductions in nitrate loadings to drinking water sources may be most valuable in places where current concentrations are close to, but not yet exceeding critical levels, whereas reductions in areas already exceeding critical levels, and where investments in treatment systems have already been made, may be less valuable.

The approach and data used for connecting changes in fertilizer application rates to beneficiary relevant outcomes represents one combination among several explored. We also considered drawing on other indicators of groundwater vulnerability to N contamination, such as DWR data on the number and density of wells (California Water Boards), USGS data on probabilities of nitrate contamination (20), and CalEnviroScreen (21). In general, the alternate approaches were deemed infeasible due inadequate information on end use (e.g. domestic use vs industrial, agricultural, etc.), being too dated, or problematic spatial correspondence. Future work may capitalize on the more resource intensive approach taken by Honeycutt et al (22).

Table SM6 -- Estimated nitrate concentrations by SWAP region.

The values are stored in the repository file “nitrate_conc_by_region.csv”.

SWAP region	Estimated nitrate concentration (mg/L)
1	5.35
2	7.14
3	4.50
4	6.61

[17] Stored values in the raster were originally normalized by the nitrate concentration to create an index. Therefore mg/L values are recoverable by multiplying by 10.

5	1.99
6	3.55
7	9.18
8	5.69
9	5.11
10	NA

Carbon flux assessment

LUCAS modeling

The land use and carbon scenario simulator (LUCAS) is an empirical model of land use change coupled with a gain–loss model of ecosystem carbon dynamics (23). We used a LUCAS model formulation that was designed to project changes in ecosystem carbon balance for the state of California under a range of climate and land-use scenarios on an annual timestep (24). Land use transitions and carbon dynamics were as described in (24), with the following modifications specific to this study:

1. Rather than using the entire statewide California LUCAS model, we ran a simplified version on a single 1-km grid cell including state classes and carbon dynamics specific to the California Central Valley Ecoregion, using only transitions relevant to the SJV (for example, we did not include the effects of wildfire or drought induced tree mortality). This simplified model also did not incorporate the various climate change scenarios described in (24), and so did not address the long-term impact of climate change on plant productivity, organic matter decomposition, and heterotrophic respiration.
2. In addition, we modified parameters associated with some of the land use transitions. Specifically:
    a. We added a “Fallow Plowing” transition, which occurred every year in the “Retired Unrestored” classification and three out of five years in the “Intermittent annuals” classification (see below). Fallow plowing represented a soil disturbance and resulted in a flux of soil carbon directly to the atmosphere in the amount of 10% of the total soil organic carbon pool, the same soil disturbance specified in the agricultural expansion transition type (24).
    b. Grasslands restored from agricultural land were assumed to consist of native perennial grasses, rather than the non-native annual grasses in the “Grassland” state class of the original California LUCAS model (24). We therefore modified restored grassland carbon dynamics so that 48% of live biomass carbon was transferred to the litter pool on an annual basis rather than 100% live biomass turnover.
    c. Sleeter et al assume specific initial soil carbon values as a function of land use history, ranging from 45.3 tC/ha for perennial (orchard & vineyard) cropland and 96.5 tC/ha for annual cropland. These are approximations, and applying this assumption without modification in our setting would confound the impact of initial soil carbon and transition type. We therefore chose to run scenarios with uniform initial soil carbon at high and low

values that encompass those used in Ref (24) (40 and 100 tons/ha respectively), and assess the sensitivity of the narrative to those assumptions. This approach recognizes that in an evolving landscape, the initial land use class for any particular snapshot in time may not be a strong predictor of the soil carbon and that the reality will be a mix of values within our sample extremes.

The default parameters for the LUCAS model are previously archived,[18] and the parameter files capturing modifications for this study are included.

Transition Scenarios

We simulated sixteen transition scenarios representing combinations of three starting classes of agriculture and five end-states, including restoration under two different vegetation types.

1) Continuous annuals (no change)
2) Intermittent annuals (no change)
3) Continuous perennials (no change)
4) Continuous annuals → Perennials
5) Continuous annuals → RetiredUnrestored
6) Continuous annuals → Grassland
7) Continuous annuals → Shrubland
8) Intermittent annuals → Perennials
9) Intermittent annuals → RetiredUnrestored
10) Intermittent annuals → Grassland
11) Intermittent annuals → Shrubland
12) Perennials → Continuous annuals
13) Perennials → Intermittent annuals
14) Perennials → RetiredUnrestored
15) Perennials → Grassland
16) Perennials → Shrubland

Model scenarios that include a transition were simulated under the original land cover for 50 years, experience a transition in the 51st modeled year, and then are run for an additional 149 years. Transitions to retirement are modeled in two ways: One with discing and one without, and all transitions are modeled with two different starting soil carbon values.

LUCAS output transformation and accounting

Each simulation of a land cover trajectory produces a time series of soil carbon stocks *S*, from which year-specific fluxes *F* can be derived, and indexed relative to a transition year 0.[19]

$$F_t = S_t - S_{t-1}\,;\ t \in \{-49,\ ...,\ 150\}$$

[18] https://www.sciencebase.gov/catalog/item/5cd08011e4b09b8c0b79a3dd

[19] STSIM output that is processed in our analysis uses the years 2000 to 2200 as year indices, but these are essentially arbitrary remnants of simulations within larger projects and do not have meaning in our context.

The net impact (net flux *NF*) that a transition (per unit area) has in a given year is identified by subtracting the flux value for the un-transitioned landscape (represented by *a*) from the transitioned landscape (represented by *b*).

$$NF_t^{a,b} = F_t^b - F_t^a$$

For carbon accounting, our main results use a time horizon of 80 years, corresponding to transitions beginning in 2021 and considering impacts to 2100. However, not all land is assumed to transition in 2020, but is rather staggered over time, discussed next.

Accounting for transitions over time[20]

If all land undergoing a specific transition did so in year zero, the GHG impact of the transition would just be the sum of net impacts from the initial time to the time horizon T (as mentioned above, our default time horizon is 80 years):

$$\sum_{t=0}^{T} NF_t^{a,b}$$

However, if some land transitions later, the expression becomes more complicated, and we must also truncate to account for the time horizon. For a unit of land converted each year over twenty years from 2021 to 2040, the contributions are staggered by year, and the total annual contribution is the sum of columns in the table below. If we assume an even distribution of transitions over the 20 year implementation period, then the average per-area effect of a transition in a given year will be those column totals divided by the length of the implementation period.

Calendar year → Transition year ↓	2021	2022	2023	2024	...	2100
2021	$NF_0^{a,b}$	$NF_1^{a,b}$	$NF_2^{a,b}$	$NF_3^{a,b}$	...	$NF_{79}^{a,b}$
2022		$NF_0^{a,b}$	$NF_1^{a,b}$	$NF_2^{a,b}$	...	$NF_{78}^{a,b}$
2023			$NF_0^{a,b}$	$NF_1^{a,b}$	...	$NF_{78}^{a,b}$
...				...	...	...
2040					...	$NF_{60}^{a,b}$

The above expressions can be used to create a multiplier table to translate the area associated with each conversion in an SLRR scenario. That is, for a given scenario, we take the acreage associated with a particular transition and multiply it by the aggregated values to estimate total GHG impact within a specific

[20] A small extension of the code allows one to apply a time-invariant social cost of carbon for benchmarking purposes, though we do not include these results in the manuscript. Under the simplifying assumption of a constant social cost of carbon, we can multiply the above values by the SCC and adjust for discounting to estimate a present value associated with the difference in emissions. More refined estimates could allow for an SCC that depends on the calendar year of emission.

accounting time horizon and land transition timeline. It can also be used to provide a multiplier for monetary impact given a particular discount rate and SCC trajectory. This created in the file carbon_staggering.R

Propagating unit impacts to landscape changes

Our analysis of carbon impacts focuses on the difference between a BAU future with SGMA and a future with SGMA implemented with retirement and restoration. This illuminates the specific value of strategic restoration approach, and also reduces the role of tenuous assumptions associated with the evolution of the landscape in other dimensions. Therefore, we estimate transitions between the present day and future landscapes, but then focus on their *difference*, so that any error in parts of the landscape unaffected by restoration is netted out.

For simplicity and due to lack of readily available data, we assume the present day landscape experiences the same within-region cropping intensity levels as the BAU landscape (a strong assumption), and that a transition to (or from) intermittently cultivated annuals can be modeled as transition to (or from) a weighted combination of continuously cultivated annuals and continuously fallowed land. For example, if a region has a cropping intensity of .75 and experiences a transition of 100 hectares from annual to perennial, that would be modeled as 75 hectares transitioning from continuously cultivated annual to perennial, and 25 hectares transitioning from continuously fallowed to perennial. Additional rules for comparing landscapes are specified with extensive comments within the file GHG_transition_propagation.R.

Results: Supplementary tables, figures, and analysis

This section contains additional results referenced in the main manuscript, as well as intermediate outputs and sensitivity analysis.

Table SR1 -- Land use and production statistics with and without SGMA

Column names:

1. #: Region number
2. Verbatim
3. Exist Footprint (ha): Present/No-SGMA footprint (ha)
4. Perm Ret (ha) : Permanent retirement (ha)
5. Perm Ret (%): Permanent retirement (%)
6. Verbatim
7. No-SGMA prod (ha): Future-without-SGMA average annual production (ha)
8. BAU prod (ha): BAU average annual production (ha)
9. SGMA prod impact (ha): SGMA production impact (ha)
10. SGMA prod impact (%): SGMA production impact (%)
11. BAU ave ann footprint (ha): BAU average annual footprint (ha)
12. BAU ave ann temp fallowing (ha): BAU average annual temporary fallowing (ha)
13. BAU ave ann temp fallowing (%): BAU average annual temporary fallowing (%)

R#	Region Name	Exist Foot-Print (ha)	Perm ret (ha)	Perm ret (%)	BAU foot-print (ha)	No-SGMA prod (ha)	BAU prod (ha)	SGMA prod impact (ha)	SGMA prod impact (%)	BAU ave ann footprint (ha)	BAU ave ann temp fallowing(ha)	BAU ave ann temp fallowing (%)
R1	Modesto	54799	0	0	54799	50778	50778	0	0	48629	6170	11.3
R2	Turlock	99424	0	0	99424	89455	89223	232	0.3	81832	17592	17.7
R3	Merced-Chowchilla-Madera	257930	21677	8.4	236253	219213	190396	28818	13.1	186241	50012	21.2
R4	Delta-Mendota	204551	3560	1.7	200991	153019	149061	3958	2.6	145750	55241	27.5
R5	Westside	236023	19358	8.2	216664	172971	148205	24765	14.3	148085	68579	31.7
R6	Kings-Tulare Lake	491442	8286	1.7	483156	364493	340543	23949	6.6	327396	155760	32.2
R7	Kaweah	139159	4864	3.5	134295	182640	170744	11896	6.5	127792	6503	4.8
R8	Tule	160577	8516	5.3	152061	207038	191760	15278	7.4	152061	0	0
R9	Kern	416308	18930	4.5	397378	293728	244538	49190	16.7	243462	153915	38.7
R10	Pleasant Valley-Kettleman Plain	33359	372	1.1	32987	8390	6329	2061	24.6	6329	26658	80.8
	SJV overall	2093570	85563	4.1	2008007	1741726	1581577	160148	9.2	1467577	540430	26.9

“Footprint” and “temporary fallowing” refer to physical land use footprint (eg, one hectare planted twice in a year has a footprint of one hectare), while “production” and “planted area” refer to gross quantities that include sums for multi-cropping (one hectare planted twice counts as two hectares). Production statistics are direct SWAP outputs, while land use statistics involve combining SWAP outputs with the initial land use map and land use change assumptions. As noted in “Estimating temporary fallowing” methods section, these assumptions result in a (slightly) negative temporary fallowing for Tule (R8), which reflects a cropping intensity greater than 1, and is therefore reported as zero temporary fallowing. BAU temporary fallowing statistics are calculated relative to the BAU footprint, not the present/no-SGMA footprint.

Table SR2 -- SJV-wide restoration statistics by land use change scenario

LUC scenario main driver	Active retirement and restoration (ha)	BAU retired in solution (ha)	Total restored area (ha)	BAU retired not in solution (ha)
even	11142	7836	18978	77727
land asset	10635	8177	18811	77387

land impairment	11369	7444	18813	78119
groundwater	11096	7996	19092	77567
surface water	9686	9024	18710	76539

Table SR3 -- Nitrogen co-benefits supplemental results

Baseline nitrate concentrations, loadings, and impact of active retirement due to strategic restoration. Median, lower, and upper refer to the range across the five scenarios. Zeros exist where no active retirement was identified as part of the restoration solution within a given region and land use change scenario. The Maximum Contaminant Level under California code is 10 mg/L, but it is important to be aware that a region-level average below that value does not imply every area within the region meets the standard.

Region	Baseline Nitrate Level (mg/L)	BAU N loading (t)	Median impact (t)	Lower impact (t)	Upper impact (t)
1	5.35	6291.85	0.00	0.00	0.00
2	7.14	11810.41	0.00	0.00	0.00
3	4.50	22072.57	1.84	1.80	4.31
4	6.61	17095.29	0.00	0.00	0.00
5	1.99	20256.21	0.00	0.00	1.14
6	3.55	33397.75	6.66	1.02	12.48
7	9.18	18768.12	0.00	0.00	0.00
8	5.69	21547.50	0.00	0.00	5.82
9	5.11	28604.84	241.65	191.72	243.86
10	NA	961.77	49.39	42.11	51.79

Soil carbon trajectories:

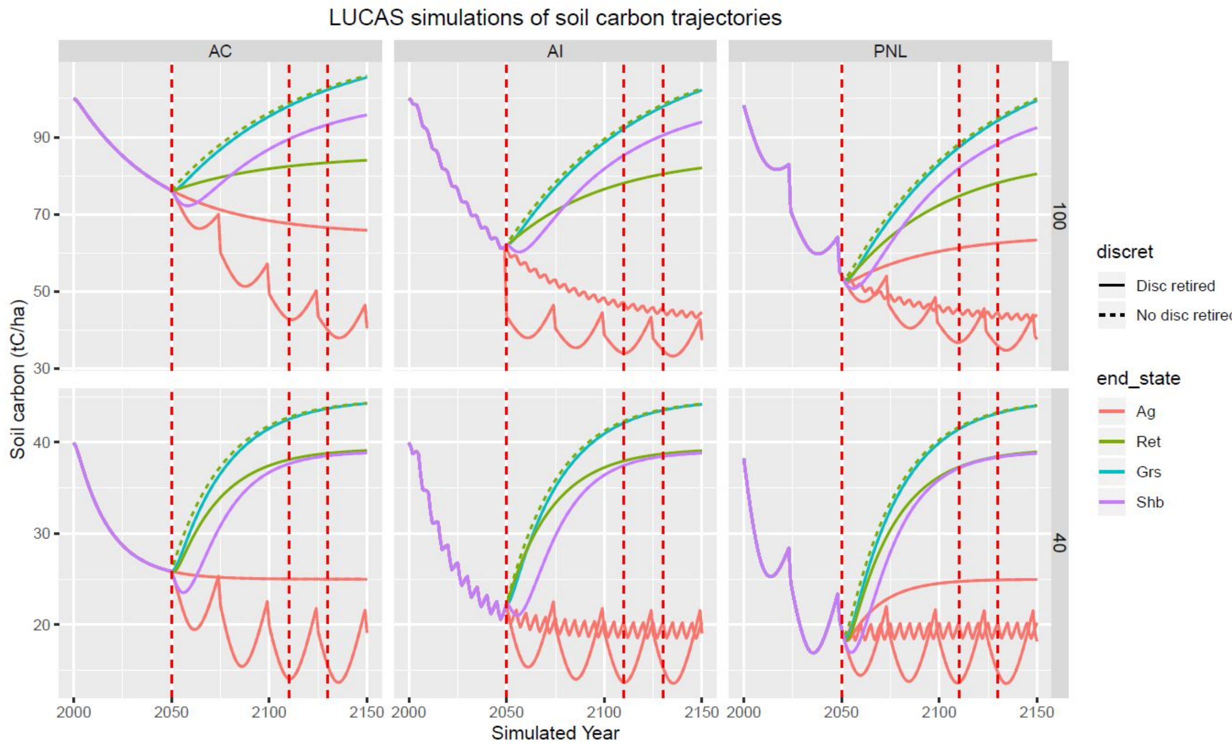

Figure SR1 -- Modeled soil carbon trajectories for all transitions.

The upper row is for soil carbon initialized to 100 tC/ha, and the lower for 40 tC/ha. The columns correspond to initial states being annual continuous, annual intermittent, and perennial. End states are color coded for Retirement, (restored) Grassland, and (restored) Shrubland. Post transition ag state is clear from the profile: Smooth is annual continuous, long-period is perennials, and short period is intermittent annuals. The x-axis labels are accounting years used within STSIM and do not correspond to any particular year. The transition is modeled at the leftmost dashed line, while the right two dashed lines bound the time horizon used for the analysis.

Uncertainty assessment for LUC and optimization

An in-depth spatial sensitivity analysis was beyond the scope of this study, however we did conduct 40 additional formal optimization runs to assess the degree to which the core findings were sensitive to key parameters and problem formulations, as well as numerous *ad hoc* explorations in the process of refining optimizer parameters. Both are described below.

Parametric sensitivity explorations

The analysis involved a one-way sensitivity assessment relative to the default values presented in the main body. One-at-a-time approaches are less preferred over designs that more thoroughly explore the space, but alternate designs are both more computationally intensive and a significantly heavier lift to parse in the spatial context. We are currently exploring future work that may examine these and other sensitivities in more detail as this framework is translated into an operational guidance within the SJV.

To assess robustness of the spatial pattern of high priority restoration areas, our primary diagnostic was the frequency map for restoration (analogous to Figure 6 in main manuscript). As with all the analysis, each varied parameter was assessed for each of the five different downscaling scenarios.

Table SR4: Parametric sensitivity explorations

Parameter	**Reference**	**Lower**	**Upper**
Species specific high-quality area target	25000 acres	75% of reference	125% of reference
Threshold for high quality habitat applied to logistic SDM output	.9	.8	NA
Threshold fraction of natural lands for a pixel to qualify as “natural” for boundary penalty	.75	.65	.85
Fraction of cost for securing BAU-retired land compared to unretired land under annuals	.1	.05	.2

With the exception of the threshold for high quality habitat, the hotspots for restoration were generally consistent across each cluster of sensitivity runs. Reducing the threshold for high quality habitat causes a consolidation of the northern Coalinga Nose (R10) component and the Western Kern (R9) component into a single area in Region 10.

Additional sensitivities considered for land use change and optimization

Optimistic SGMA scenario

The SWAP modeling included a lower-impact SGMA scenario associated with significant investments in new water supplies. As mentioned in the SWAP modeling section, stakeholders did not find the level of fallowing and retirement realistic and so we do not focus our analysis on that scenario. However, we did conduct some optimization runs after propagating the lower-impact SGMA scenario through our workflow, and these indicate that patterns for restoration remain primarily the same, except that lower retirement and overall higher returns in Tule (R8) cause a general (but not complete) shift in recommended restoration from Tule to Kern (R9).

Resolution of analysis

Early runs also examined the effect of running the land use change and optimization at an alternate resolution of 270 meter pixels rather than 1080 meter pixels. While there were some differences particularly related to smaller clusters, we determined that, for the purposes of this study, overall patterns were not sufficiently different to merit the coding effort required to preserve the generality of the workflow for multiple resolutions. Moreover, the scale of analysis has an implementation interpretation related to the minimum size of land worth bringing under restoration, and we deemed the one kilometer scale to be desirable for ensuring reasonably sized restored areas, though it is not a crisply defined level and we are exploring future work to test different resolutions more thoroughly.

Boundary penalty

Early runs also explored variations in the boundary penalty. In the chosen least cost problem framing, the boundary penalty indicates the willingness to pay for a more tightly clustered set of natural areas. The boundary penalty used for this analysis ($100 per meter) was chosen after examining results produced across a range of values and making a subjective assessment regarding the degree of clustering and resulting patch sizes. A more extensive analysis could identify a boundary penalty that corresponds to key patch sizes, however this was beyond the scope of this study and also may constitute an exercise in false precision, given that this study intends to identify the approximate pallette within with lands may be secured for restoration, rather than exact reserve design.

References

1. R. E. Howitt, J. Medellín-Azuara, D. MacEwan, J. R. Lund, Calibrating disaggregate economic models of agricultural production and water management. *Environ. Model. Softw.* **38**, 244–258 (2012).
2. P. Mérel, R. Howitt, Theory and Application of Positive Mathematical Programming in Agriculture and the Environment. *Annu. Rev. Resour. Econ.* **6**, 451–470 (2014).
3. E. Jr. Thompson, D. Pearce, “San Joaquin Land and Water Strategy” (American Farmland Trust, 2018).
4. T. E. Grantham, J. H. Viers, 100 years of California’s water rights system: patterns, trends and uncertainty. *Environ. Res. Lett.* **9**, 084012 (2014).
5. California Department of Water Resources (CADWR), *Groundwater basin prioritization: final CASGEM basin prioritization results—June 2014* (2014) (May 31, 2018).
6. J. E. Ayars, C. J. Phene, “Improving Pomegranate Fertigation and Nitrogen Use Efficiency with Drip Irrigation Systems. Project Planning Document” (USDA-ARS, 2011).
7. R. Y. Evans, L. Dodge, J. Newman, Nutrient Management in Nursery and Floriculture (2007).
8. D. Munier, *et al.*, “Sample costs to produce safflower” (Sacramento Valley, University of California Cooperative Extension, 2011).
9. E. J. Perry, K. M. Klonsky, J. J. Shimada, “Sample costs to establish and produce boysenberries” (San Joaquin Valley-North, University of California Cooperative Extension, 2002).
10. T. P. Tomich, S. B. Brodt, R. A. Dahlgren, K. M. Scow, Eds., *The California Nitrogen Assessment: Challenges and solutions for people, agriculture, and the environment* (University of California Press, 2016).
11. J. Abbadi, J. Gerendás, Nitrogen Use Efficiency of Safflower as Compared to Sunflower. *J. Plant Nutr.* **32**, 929–945 (2009).
12. C. Carranca, G. Brunetto, M. Tagliavini, *Nitrogen Nutrition of Fruit Trees to Reconcile Productivity*

and Environmental Concerns (2018).
13. T. A. Doerge, K. H. Pritchard, J. W. Pier, P. Fernandez, T. W. McCreary, *Nitrogen Utilization Efficiency in Melons Using Soluble and Slow Release Fertilizers*, N. F. Oebker, M. Bantlin, Eds. (1991).
14. R. Hermanson, W. Pan, C. Perillo, R. Stevens, C. Stockle, Nitrogen use by crops and the fate of nitrogen in the soil and vadose zone (2000).
15. J. P. Jones, T. C. Tucker, Effect of Nitrogen Fertilizer on Yield, Nitrogen Content, and Yield Components of Safflower. *Agron. J.* **60**, 633–634 (1968).
16. S. F. Ledgard, G. S. Smith, M. S. Sprosen, Fate of 15N-labelled nitrogen fertilizer applied to kiwifruit (Actinidia deliciosa) vines. *Plant Soil* **147**, 49–57 (1992).
17. B. C. Strik, D. R. Bryla, Uptake and Partitioning of Nutrients in Blackberry and Raspberry and Evaluating Plant Nutrient Status for Accurate Assessment of Fertilizer Requirements. *HortTechnology Hortte* **25**, 452–459 (2015).
18. A. Sorgonà, M. Abenavoli, “Nitrogen in Citrus: Signal, Nutrient, and Use Efficiency” in *Advances in Citrus Nutrition*, (2012), pp. 231–244.
19. T. Harter, *et al.*, “Nitrogen Fertilizer Loading to Groundwater in the Central Valley” (University of California, Davis, 2017).
20. B. T. Nolan, K. J. Hitt, B. C. Ruddy, Probability of nitrate contamination of recently recharged groundwaters in the conterminous United States. *Environ. Sci. Technol.* **36**, 2138–2145 (2002).
21. J. Faust, *et al.*, “CalEnviroScreen 3.0” (California Environmental Protection Agency and Office of Environmental Health Hazard Assessment, 2017).
22. K. Honeycutt, H. Canada, M. Jenkins, J. Lund, “Alternative Water Supply Options for Nitrate Contamination” (State Water Resources Control Board, 2012).
23. C. J. Daniel, B. M. Sleeter, L. Frid, M.-J. Fortin, Integrating continuous stocks and flows into state-and-transition simulation models of landscape change. *Methods Ecol. Evol.* (2018) https:/doi.org/10.1111/2041-210X.12952 (February 28, 2018).
24. B. M. Sleeter, *et al.*, Effects of 21st-century climate, land use, and disturbances on ecosystem carbon balance in California. *Glob. Change Biol.* **25**, 3334–3353 (2019).