1995a). SWAP covers over 93% of irrigated agriculture in California, most of which is in the Central Valley, and calibrates exactly to an observed base year of land use and input allocation data through use of exogenous elasticities and assumed profit-maximization behavior by farmers. This paper, (i) documents SWAP and motivates application to other regions, (ii) discusses SWAP construction and emphasizes the sequential calibration diagnostic checks used in the model, (iii) extends the applied PMP literature with more flexible production and cost functions, and (iv) links the SWAP model to the infrastructure of a hydro-economic network model for water supply in California (CALVIN). We conclude with an empirical example and estimate the value of water markets for California's San Joaquin Valley.

The next section highlights the importance of micro-level policy analysis, in various geographic regions, with models similar to SWAP. We place SWAP in the context of the existing literature of optimization models and PMP. In the subsequent section we construct SWAP with particular emphasis on the sequential calibration routine and improvements over previous PMP models. The calibration routine has six steps with model consistency checks at each stage. Improvements over similar PMP models include Constant Elasticity of Substitution (CES) production functions, exponential PMP land cost functions, and endogenous crop prices. Finally, we demonstrate an application of SWAP for evaluation of water markets in the San Joaquin Valley. We conclude the paper with a discussion of extensions, limitations, and future work on SWAP.

1.1. Micro-level analysis of agricultural policies

In the U.S. and other agricultural economies the demand for micro-level analysis of agricultural policies that reflect the effects on local agricultural and environmental resources is growing for several reasons. National agricultural policies are increasingly driven or constrained by environmental criteria. Furthermore, there are an increasing number of regional (state) level policies that proscribe the use of agricultural inputs (land, water, labor and supplies) and resources. The era of unfettered commodity price support programs whose impact could be measured by aggregate financial or physical outcomes is waning, as are the aggregate demand and supply methods used to measure such outcomes. Finally, the complex physical and economic interaction between the environment and agricultural policies is difficult to accurately capture using standard econometric techniques based on aggregate data.

Calibrated optimization models for micro-level analysis, such as SWAP, focus on spatially heterogeneous commodity, resource, and input specific policies. Instead of using data from the outcome of economic optimization to estimate aggregate elasticities, calibrated optimization models use prior estimates of elasticities of demand, supply, and substitution coupled with observed micro-input data on regional production to calibrate the model. In the SWAP model we additionally assume that profit-maximizing behavior and short run equilibrium conditions led to the observed base year resource allocation. Since these models use an explicit primal specification of agricultural production, they can model policies defined in terms of physical resource limits rather than financial outcomes.

1.2. Optimization models and Positive Mathematical Programming

Moore and Hedges (1963) first introduced models of irrigated agriculture as a way to estimate irrigation water demand. They, and later studies, used mathematical (typically linear) programming models to estimate irrigation water demand elasticities. Gardner (1983) reviewed studies on irrigation water demand, with emphasis on California, completed during the 1960s and 1970s. This literature has since evolved to focus on large-scale regional optimization models. Today optimization models are used to analyze water demand and agricultural–environmental policies, since these models work better with a multitude of resource constraints and complex interactions between agriculture and the environment (Griffin, 2006).

A major problem that initially plagued optimization models was a tendency to overspecialize in crop production (Howitt, 1995a). In response, the 1980's saw the first models based on the technique of Positive Mathematical Programming (PMP). PMP is a deductive approach to simulating the effects of policy changes on cropping patterns at the extensive and intensive margins. The term "positive" implies the use of observed data as part of the model calibration process. PMP has several advantages over traditional optimization models. First, the PMP cost function calibrates the model exactly to observed values of production output and factor usage. Second, PMP adds flexibility to the profit function by relaxing the restrictive linear cost assumption. A third advantage is that PMP does not require large datasets. Heckelei and Britz (2005) note that PMP models can be viewed as a bridge between econometric models, with substantial data requirements, and more limited traditional optimization models. Finally, programming models including the subset of PMP models such as SWAP are more responsive to policy changes than statistical (inductive) models of agricultural production (Scheierling et al., 2006).

Calibration of production models by PMP has been reviewed extensively in the literature and variations on the base method have been developed. Buysse et al. (2007) and Heckelei and Wolff (2003) argue that shadow values from calibration and resource constraints are an arbitrary source of information for model calibration. Subsequent research suggests the use of exogenous information such as land rents instead of shadow values (Heckelei and Britz, 2005; Kanellopoulos et al., 2010). Heckelei and Britz (2005) and Paris and Howitt (1998) propose a generalized maximum entropy (GME) formulation to estimate resource and calibration constraint shadow values. However, the GME procedure has seen little use in applied research. Merel and Bucaram (2010) and Merel et al. (2011) propose calibration against exogenous, and potentially regionally-disaggregate, supply elasticity estimates.

Research on linked hydrologic and economic models has evolved parallel to research on PMP with a focus on improved policy simulations and analysis. Economic models typically omit a hydrologic representation and hydrologic models lack the ability to economically allocate water. Hybrid hydrologic–economic models can be holistic (one model) or compartmental (sequential iteration between different models) (Cai, 2008; Braat and vanLierop, 1987). Compartmental hydrologic–economic models are frequently a hydrologic model linked with an economic model calibrated by PMP. Gomann et al. (2005) link the RAUMIS economic model, calibrated using PMP, to GROWA98 and WEKU hydrologic models to model the effects of Nitrogen tax relative to a quota on dairy herds to increase water quality in Germany. In an example of work in California, Quinn et al. (2004) adopt a compartmental approach and develop the PMP APSIDE economic model which is linked to the CALSIM II water model. They also include climate simulations, in a third model, to evaluate climate change impacts in California. vanWalsum et al. (2008) introduce the bio-economic model Waterwise which is linked to the DRAM PMP model. They use the model to evaluate European Union water quality policies in the Netherlands.

Despite the many papers employing PMP models to infer economic values for water and environmental resources, we cannot find any publication that focuses on the calibration procedure for PMP economic models and formal diagnostic tests for each calibration stage. The calibration-diagnostic iterative procedure applies to standalone economic models and linked hydrologic–economic models in diverse geographic regions.

2. The Statewide Agricultural Production Model (SWAP)

2.1. SWAP modeling framework

A model is, by definition, a simplified representation of a real system. In the process of abstracting and simplifying a real system, a model loses some information; thus even with theoretically consistent structure it is unlikely that a model will calibrate closely to observed (base year) data. The problem is well documented in agricultural production modeling (Hazell and Norton, 1986). One solution is to use observed farmer behavior, in the form of observed land use patterns, and additional exogenous information to calibrate parameters of the structural model that exactly reproduce observed base-year conditions. Positive Mathematical Programming is a common calibration method for structural agricultural production models.

The SWAP model is a regional model of irrigated agriculture in California, calibrated using PMP. PMP can derive model parameters so that first-order conditions for economic optimization are satisfied at an observed base year of input and output data. This is accomplished by assumed profit maximizing behavior by farmers and a non-linear objective function. SWAP offers three key improvements over traditional PMP models. First, SWAP includes regional exponential PMP land cost functions, which corrects the inability of previous models, with quadratic functions, to handle large policy shocks. Second, SWAP includes regional Constant Elasticity of Substitution (CES) crop production functions which allow limited substitution between inputs. Leontief production functions were common in most previous models. Finally, regional crop prices are endogenously determined based on a statewide demand function.

SWAP was originally developed to be the agricultural economic component for the CALVIN model of the California water system (Draper et al., 2003). It has subsequently been used in a wide range of policy analyses in California. SWAP has been used to estimate economic losses due to salinity in the Central Valley (Howitt et al., 2009), economic losses to agriculture due to alternative conveyance in the Sacramento-San Joaquin Delta (Appendix to Lund et al., 2007), economic losses to agriculture and confined animal operations in California's Southern Central Valley (Medellin-Azuara et al., 2008), and economic effects of water shortage on Central Valley agriculture (Howitt et al., 2011). The model has also been linked to agronomic yield models in order to estimate effects of climate change on irrigated agriculture in California (Medellin-Azuara et al., 2012). Variations of SWAP also have been applied in other regions such as the US-Mexico border basins (Howitt and Medellín-Azuara, 2008; Medellin-Azuara et al., 2009). The model is used for policy analysis by the California Department of Water Resources (DWR, 2009) and the United States (U.S.) Department of the Interior (Interior), Bureau of Reclamation (Reclamation, 2011).

SWAP is defined over homogenous agricultural regions and assumes that farmers maximize profits subject to resource, technical, and market constraints. Farmers sell and buy in competitive markets where any one farmer cannot affect the price of any commodity. The model selects crops, water supplies, and other inputs that maximize profit subject to constraints on water and land, and subject to economic conditions regarding prices, yields, and costs. The model incorporates water supplies from state and federal projects, local water supplies, and groundwater. As conditions change within a SWAP region (e.g. the quantity of available project water supply increases or the cost of groundwater pumping increases) the model optimizes production at both the extensive and intensive margins by adjusting the crop mix, water sources and quantities used, and other inputs. It will also fallow land in response to resource conditions.

The SWAP model is written in GAMS (General Algebraic Modeling System) and solved using the non-linear solver CONOPT-3. The objective is to maximize the sum of producer (regional profits) and consumer surplus.

2.2. Model development and calibration

Development of the SWAP model is divided into calibration and policy analysis phases. Calibration is analogous to parameter estimation in econometric models or calibration in Computable General Equilibrium (CGE) models. Policy analysis estimates the effects of changing prices, costs, resources, or institutions given the calibrated parameter values.

We detail the calibration procedure for SWAP and emphasize model improvements and diagnostic checks in the process. The calibration procedure for SWAP reflects most of the ten steps discussed in Jakeman et al. (2006) with particular emphasis on sequential calibration and a parallel set of diagnostic tests to check model performance. Stepwise model development procedures have been applied for many modeling problems, including neural networks (Piuleac et al., 2010), and computational fluid dynamics (Blocken and Gualtieri, 2012). The stepwise tests specified in Fig. 1 are ordered in a logical sequence. For example, the first test for positive net returns is a necessary condition for an optimal solution in the calibrated linear program. Likewise, the equality of the input marginal value products to their opportunity costs is a necessary

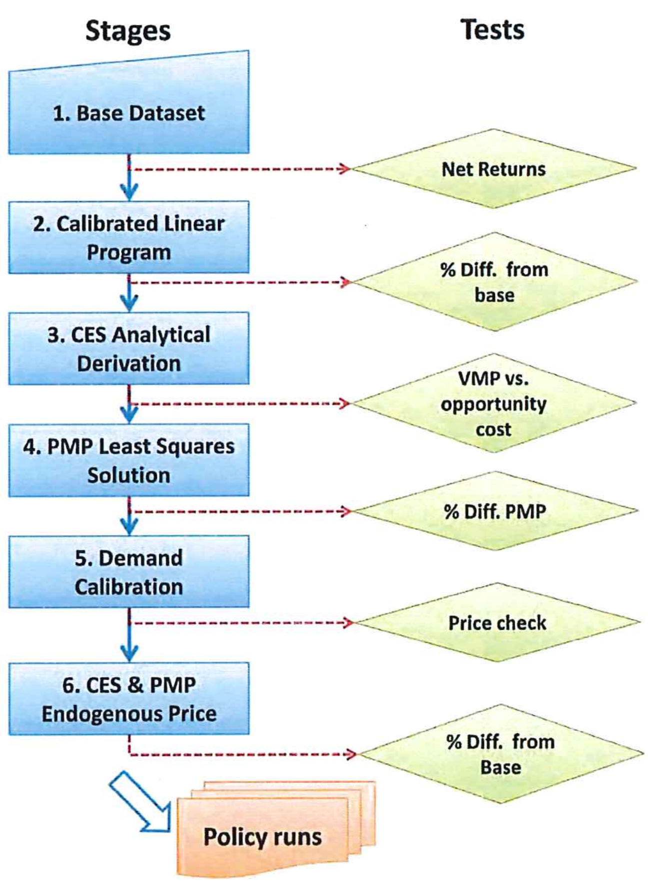

Fig. 1. SWAP calibration stages and tests.

condition for optimal input calibration in the nonlinear CES model. The sequential tests defined in Fig. 1 are a blueprint for model validation and identification of potential problems.

The calibration phase of the SWAP model uses a sequential six-step process outlined in Fig. 1. The six steps are (i) assemble input, output and elasticity data, (ii) solve a linear program subject to fixed resource and calibration constraints, (iii) derive the CES production function parameters using input opportunity costs from step two, (iv) estimate the crop and region-specific PMP cost functions using a least squares method, (v) calibrate the aggregate demand functions and regional adjustment costs using prior demand elasticity estimates, and (vi) optimize and simulate the calibrated SWAP model which includes tests for adequate calibration in terms of input and output prices and quantities.

Model calibration data should be representative of "normal" production conditions in the relevant region. We take 2005 as the base year in the SWAP model because it represents the most recent data available for an average water and price year in California. The model calibrates to the base year in terms of the following parameters: crop output quantities, output prices, input quantities, input value marginal products, variable costs, and imputed costs to fixed inputs.

2.2.1. Step I: data assembly

The level of spatial aggregation is important for defining the scope and method of analysis. Disaggregated production models typically require more data but tend to be effective in policy analysis in rural economies (Taylor et al., 2005). When agricultural production is homogeneous and production conditions are relatively stable, there is less information gained from disaggregation. SWAP aggregates agricultural production data to the level of representative regions. The SWAP regions are based on the California Department of Water Resources (DWR) Detailed Analysis Units (DAU). Each SWAP region is composed of one or more DAU with homogenous microclimate, water availability, and production conditions. This scale is more suitable for statewide hydro-economic models that require marginal economic values of water for competing agricultural and urban demand locations (Draper et al., 2003). The SWAP model has 27 base regions in the Central Valley plus the Central Coast, the Colorado River region that includes Coachella, Palo Verde and the Imperial Valley and San Diego, Santa Ana and Ventura, and the South Coast. The model has a total of 37 agricultural regions, only 27 regions in the Central Valley are considered for the analysis in this paper. Fig. 2 shows California agricultural area covered in SWAP.

We aggregate crops into 20 representative crop groups. A single crop group can represent several individual crops. Irrigated land use represents the area of all crops within the group, production costs and returns are represented by a single proxy crop for each group. The current 20 crop groups were defined in collaboration with DWR (DWR, 2010). For each group we choose the representative (proxy) crop based on four criteria: (i) availability of

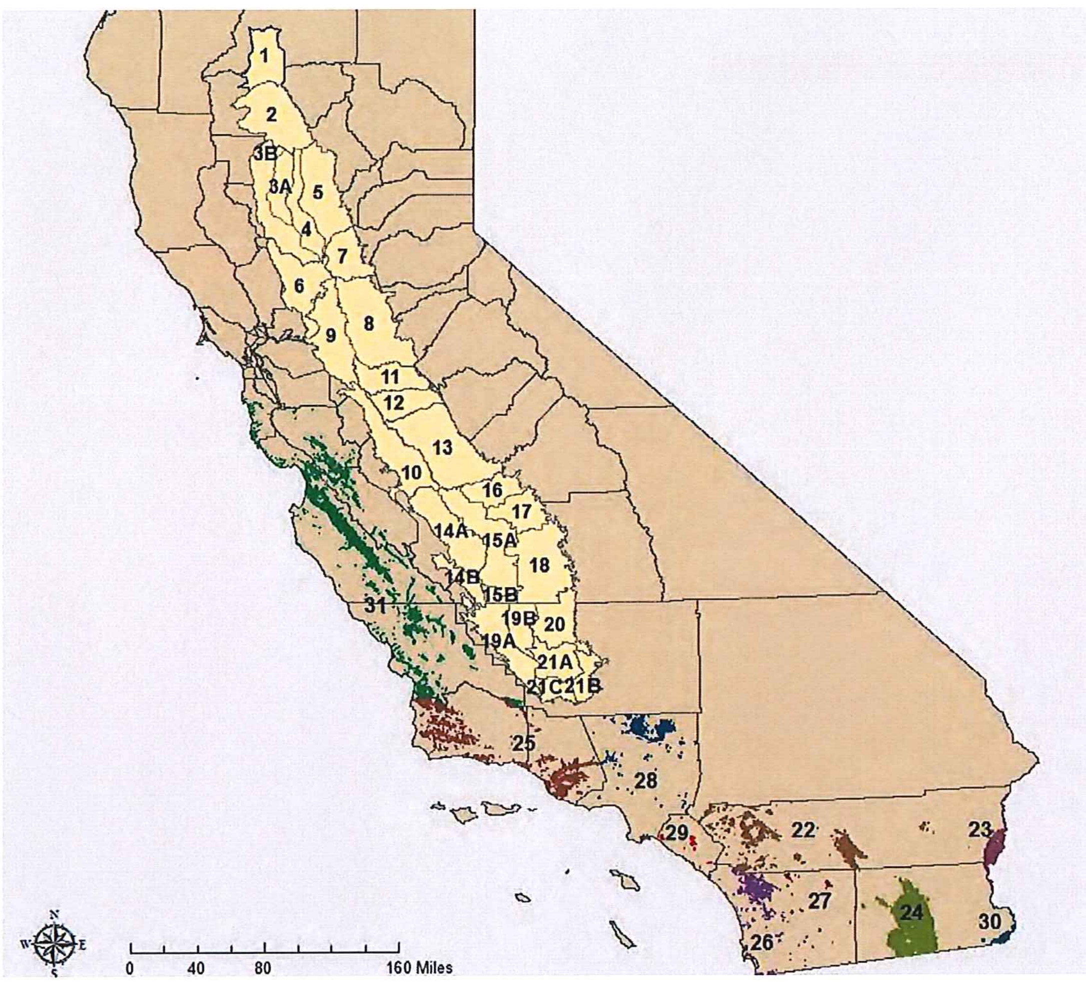

Fig. 2. SWAP region definition and coverage.

a detailed production budget, (ii) representative of the largest land use within a group, (iii) representative of water use (applied water) of all crops in the group, and (iv) having similar gross and net returns as other crops in the group. The relative importance of these criteria varies by crop. The 20 crop groups include almonds and pistachios, alfalfa, corn, cotton, cucurbits, dry beans, fresh tomatoes, grains, onions and garlic, other deciduous, other field, other truck, irrigated pasture, potatoes, processing tomatoes, rice, safflower, sugar beets, subtropical, and vines.

Variable input costs for the crop groups are derived from the regional cost and return studies from the University of California Cooperative Extension (UCCE, 2011). There are four aggregate inputs to production, (i) land, (ii), labor, (iii) water, and (iv) other supplies. All inputs except water are derived from the UCCE Budgets. Since cost budgets represent best management practices, SWAP also uses the corresponding yields from the budgets. Commodity prices for the base year in the model are from the California County Agricultural Commissioner's reports published by the U.S. Department of Agriculture (USDA, 2011).

We derive applied water per hectare (base) requirements for crops in SWAP from DWR estimates (DWR, 2010). DWR estimates are based on Detailed Analysis Units (DAU). An average of DAU's within a SWAP region is used to generate a SWAP region specific estimate of applied water per hectare for SWAP crops.

The SWAP model includes five types of surface water: State Water Project (SWP) delivery, three categories of Central Valley Project (CVP) delivery, and local surface water delivery or direct diversion (LOC). The three categories of CVP deliveries represent water service contract and include Friant Class 1 (CVP1), Friant Class 2 (CL2), and water rights settlement and exchange delivery (CVPS). CVP and SWP water costs have two components, a project charge and a district charge. The sum of these components is the region-specific cost of the individual water source.

Groundwater pumping costs are calculated as two components, the fixed cost per cubic meter based on typical well designs and costs within the region, plus the variable cost per cubic meter. The variable cost per cubic meter is O&M plus energy costs based on average total dynamic lift within the region. In our example application we consider a short run drought analysis and hold dynamic lift and groundwater pumping costs constant. Long run policy analysis may link the SWAP model to a groundwater model such as the Central Valley Hydrologic Model (CVHM) to simultaneously estimate changes in regional depth to groundwater (Reclamation, 2011).

The model calibration approach, discussed in the following section, is driven by the first order conditions and fixed resource constraints. Since the underlying objective is to maximize profits, subject to inequality constraints on the fixed inputs, each regional crop production activity must have a positive gross margin at the base calibration values. As such, the essential test at this stage is to ensure that the gross margin over variable costs is positive for those crops actually grown. If the net returns to land and management are negative after checking the data, there are several ways of addressing the problem. The simplest approach may be to use a lower bound calibration constraint in Step II to calculate the needed reduction in the land opportunity cost from the lower bound constraint shadow value. More generally, the researcher should consult extension agents and other experts to identify potential inconsistencies in the crop budgets or other input data.

2.2.2. Step II: linear calibration program

In this step we solve a linear program of farm profit maximization with calibration constraints set to observed values of land use. All other production inputs are normalized to land. The Lagrangian multipliers on the calibration and resource constraints are used in steps three and four to parameterize regional CES production functions and exponential PMP cost functions. We define sub-index g for (27) agricultural (SWAP) regions, i for (20) crop groups, j for (4) production inputs, and w for (6) individual water sources.

We solve a linear program to obtain marginal values on calibration and resource constraints. The linear program objective function is to maximize the sum of regional profits across all crops by optimizing land use $xl_{gi,land}$ and water use $watl_{gw}$. Equation (1) defines the objective function,

$$\max_{xl_{gi,land}, watl_{gw}} \Pi = \sum_{g}\sum_{i}\left(v_{gi}yld_{gi} - \sum_{j\neq water}\omega_{gij}a_{gij}\right)xl_{gi,land} - \sum_{g}\sum_{w}(watl_{gw}\varpi_{gw}), \tag{1}$$

where v_{gi} are region-specific crop prices (marginal revenue per tonne of output), yld_{gi} are the base yields for crop i in region g, ω_{gij} are input costs, ϖ_{gw} are water costs, and a_{gij} are regional Leontief coefficients defined in Equation (2). $\tilde{x}_{gij}$ represents the *observed* level of input use.

$$a_{gij} = \frac{\tilde{x}_{gij}}{\tilde{x}_{gi,land}} \tag{2}$$

Production is constrained by resource availability of binding inputs including land and water. These are treated separately in the calibration program, since regions may be binding in land, water, or both. The land resource constraints are defined as

$$\sum_{i} xl_{gi} \le b_{g,land} \quad \forall g, \tag{3}$$

where $b_{g,land}$ are region-specific land availability constraints. The water constraints are defined by region and water source,

$$\sum_{i} aw_{gi}xl_{gi} \le \sum_{w} watl_{g,w} \quad \forall g, \tag{4}$$

and

$$\sum_{w} wat_{gw} \le \sum_{w} \mathrm{watcons}_{gw} \quad \forall g \tag{5}$$

where $\mathrm{watcons}_{gw}$ are region and water source-specific constraints, and aw_{gi} are crop water requirements (applied water per hectare) and may reflect regional difference in average irrigation efficiency or consumptive use. Define λ_g^L and λ_g^W as the shadow values for Equations (3) and (4), respectively.

A calibration constraint forces the program to reproduce base year observed cropping patterns. We include a perturbation ($\varepsilon = 0.0001$) to decouple the resource and calibration constraints as detailed in Howitt (1995a),

$$x_{gi,land} \le \tilde{x}_{gi,land} + \varepsilon \quad \forall\, g, i. \tag{6}$$

We add the calibration constraint to land only, and use the shadow value of land λ_{gi}^C as the marginal price needed to calibrate optimal land allocation in Equation (6). The other inputs are calibrated by using the first order conditions for the CES production function defined later in the process.

Two tests are applied to the output of the Step II model. The first test measures any deviation in regional crop input allocation by the model. Percentage deviations in input use by crop and region of less than 1% are permissible given the small perturbations in the calibration constraints, but any input deviation greater than this implies negative gross margins, or unduly restrictive fixed input constraints. The second calibration test verifies that the number of non-zero dual values on calibration constraints plus the number of non-zero shadow values on binding resource constraints equal the number of non-zero production activities in each region. If this test does not hold, the model will not have sufficient cost information to calibrate the full set of non-zero activities as some crops should have interior solutions, but do not have calibration shadow values to derive them.

2.2.3. Step III: production function parameter calibration

In this step we sequentially derive the parameters for the Constant Returns to Scale (CRS) CES production function for each region and crop following the procedure developed in Howitt (1995b). The CES is a flexible functional form which allows for a constant rate of substitution between production inputs and nests Leontief (fixed proportions) and Cobb–Douglas (unit substitution) production technologies. Researchers use various types of quadratic functions in agricultural optimization models (Cai, 2008). The model which preceded SWAP in California, the Central Valley Production Model, modeled production along the water use-irrigation efficiency isoquant (Reclamation, 1997). SWAP improves previous methods and calibrates a CES production function for each crop and region. One key property of the CES production function is that it defines the rates at which inputs can be substituted for each other, for example, applied water used in irrigation can be partly substituted for by increased irrigation efficiency which requires additional labor and capital.

The Constant Returns to Scale (CRS) CES production functions for every region and first-order conditions for an optimum input allocation yield a sequential set of conditions to solve for the parameters of the CES. The theoretical properties may be found in Beattie and Taylor (1985). We define the CES functions as

$$y_{gi} = \tau_{gi}\left[\beta_{gi1}x_{gi1}^{\rho_i} + \beta_{gi2}x_{gi2}^{\rho_i} + \ldots + \beta_{gij}x_{gij}^{\rho_i}\right]^{v/\rho_i}, \tag{7}$$

where y_{gi} represents output of crop i in tonnes for region g, by combining aggregate inputs j. The scale parameters are (τ_{gi}) and the relative use of production factors is represented by the share parameters β_{gij}. Production factor use is given by x_{gij}. The returns to scale coefficient is v and CRS requires that the coefficient is set at 1.

The SWAP model uses a non-nested CES production function with the same elasticity of substitution between any two inputs. The SWAP model is also able to handle a nested-CES production function with two or more sub-nests and corresponding versions of the model have been developed. If data are available the substitution elasticity should be estimated. If substitution elasticities are available from existing studies those can be used. Currently there are insufficient data to estimate the elasticity of substitution, thus the value is fixed at $\sigma = 0.17$ for all inputs. We assume this value to allow for limited substitution between inputs based on experience from previous analyses.

Limited substitution between inputs is consistent with observed farmer production practices. Namely, we observe that farmers can, over a limited range, substitute among inputs in order to achieve the same level of production. Fig. 3 shows an example of a CES production surface. To show the CES function as a 3-dimensional surface two inputs (supplies and land) are held constant. The vertical axis shows total production of alfalfa in Region 15 given different combinations of water and labor which are shown on the horizontal axes. Fig. 3 illustrates two important aspects of the CES production function. First, substitution between inputs can be seen by holding production constant (the vertical axis) and sliding around the production surface. There is limited substitution between water and labor, as shown by the "sharp" corners to the production surface. Second, Fig. 3 demonstrates the ability of SWAP to model deficit (stress) irrigation by farmers or, more generally, the marginal product of a given input. Faced with a water shortage we expect that farmers may deficit-irrigate some crops. Holding labor constant and sliding along the production surface, as water is decreased production (yield) decreases as well. Additional restrictions can be imposed to incorporate exogenous agronomic data.

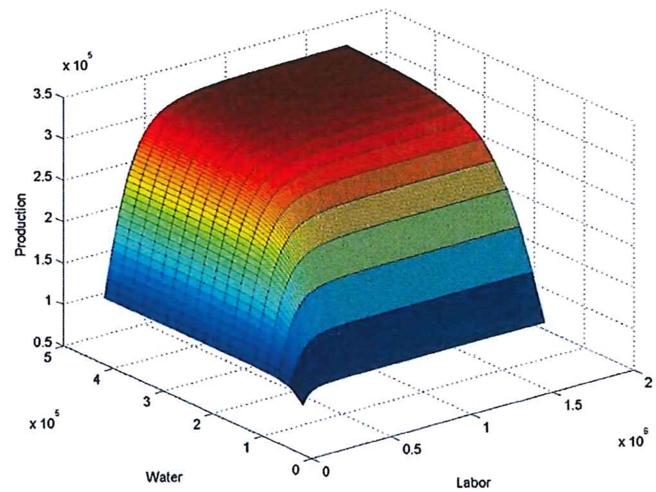

Fig. 3. Simplified CES production function surface for alfalfa in region 15.

The first order condition for optimal input allocation is that the value marginal product (output price times the marginal product) of each input for each crop and region is equal to the marginal cash cost plus opportunity cost of the input. This is equal to the base input price plus the dual value on the resource constraints, λ_g^L and λ_g^W, and, when binding, the dual value on the calibration constraint, λ_{gi}^C. The linear program in Step II will not have calibration shadow values for activities associated with the binding resource constraints. In the absence of prior estimates of the marginal productivity of these crops, we impose the assumption that marginal productivity decreases 25% over the base condition productivity and thus use 25% of the land resource shadow value as a proxy for the calibration shadow value, and adjust the other calibration values accordingly. While this is a general assumption over different regions and crops, it provides a robust method for full calibration of all the observed crops without inducing infeasibilities from more arbitrary exogenous restrictions.

Let the cost per unit of each input, inclusive of marginal cash cost and opportunity cost of input j be ω_j. To simplify notation, consider a single crop and region and normalize the price per unit output to 1. Define

$$\rho = \frac{\sigma - 1}{\sigma}, \tag{8}$$

and the corresponding farm profit maximization problem, optimizing over input use X_j, is written as,

$$\max_{X_j} \pi = \tau\left[\sum_j \beta_j x_j^{\rho}\right]^{v/\rho} - \sum_j \omega_j x_j. \tag{9}$$

Constant returns to scale requires that $v = 1$ and

$$\sum_j \beta_j = 1. \tag{10}$$

We use the restrictions imposed by constant returns to scale and take ratios of any two first order conditions to derive the familiar optimality condition that marginal rate of technical substitution equals the ratio of input costs. Let l correspond to all $j\neq 1$ and by rearranging and using the restriction in Equation (10) we can explicitly solve for the first (or any arbitrary) coefficient,

$$\beta_1 = \frac{1}{1+\frac{x_1^{(-1/\sigma)}}{\omega_1}\left(\sum_l \frac{\omega_l}{x_l^{(-1/\sigma)}}\right)}. \tag{11}$$

We use the same procedure as above for all other β_l where $l\neq 1$, thus

$$\beta_l = \frac{1}{1+\frac{x_1^{(-1/\sigma)}}{\omega_1}\left(\sum_l \frac{\omega_l}{x_l^{(-1/\sigma)}}\right)} \frac{\omega_l\, x_1^{-1/\sigma}}{\omega_1\, x_l^{-1/\sigma}}. \tag{12}$$

We calculate the scale parameter, for each region and crop, from the definition of the CES production function, evaluated at the base level. The scale parameter is

$$\tau = \frac{\left(yld/\tilde{x}_{land}\right)\cdot \tilde{x}_{land}}{\left[\sum_j \beta_j x_j^{\rho}\right]^{\upsilon/\rho_i}}. \tag{13}$$

The process generalizes to any number of regions and crops. In SWAP this process is automatically performed for all crops and regions and the production functions are fully calibrated.

2.2.3.1. Numerical scaling issues in optimization models. From the first order conditions we see that

$$\beta_l = \frac{\omega_l}{\omega_1}\frac{\beta_1 x_1^{(-1/\sigma)}}{x_l^{(-1/\sigma)}}, \tag{14}$$

for any given input l. If input costs (marginal cash cost plus opportunity cost) of two inputs are of a different order of magnitude this can cause the β_j coefficients to become unbalanced and lead to numerical issues with model calibration. Specifically, an ill-conditioned calibration routine will tend to set $\beta_l \approx 1$ and all other $\beta_j \approx 0$. In turn, the model will not calibrate with a low elasticity of substitution (large value in the exponent). This type of data issue is common with large-scale regional production models since inputs are aggregated into coarse categories. For example, other supplies have a much larger cost per unit land than labor costs for many crops causing ill-conditioned matrices that impede numerical convergence to an optimal solution.

There are many sophisticated scaling approaches but a simple solution used in SWAP is to numerically scale input costs into units of the same order of magnitude. We use land costs as the reference scale and convert input costs, except for land, into land units. We calculate the ratio of input use to total hectares, for each crop and region, and normalize the costs of production into the corresponding unit. This scaling is used throughout the SWAP program. At the end of the program we use a de-scaling routine which simply reverses this process to convert input use and costs back into standard units.

2.2.4. Step IV: estimating an exponential PMP cost function

The SWAP model posits that farmers cultivate the best land first for any given crop so additional land put into production will be of lower quality. The effect will vary over space and will depend on several additional factors including management skills, field-specific physical capital, and the dynamic effects of crop rotation. In general, additional land into production requires a higher cost to prepare and cultivate. We combine this unobservable (directly) information with average production costs to calibrate exponential land cost functions in the model.

PMP land cost functions are calibrated using information from acreage response elasticities and shadow values (implied values) on calibration constraints. Merel and Bucaram (2010) derive conditions for the exact calibration to elasticities for the Leontief and CES model with a quadratic PMP cost function. They show that the approach used here can be defined as myopic calibration, since it does not account for the effect of crop interdependency on the marginal elasticity. However they do show that under so-called "number of crops" and "dominant response" conditions, the myopic approach can be an adequate approximation. With 20 representative crops, the SWAP model is likely to satisfy both conditions, though we have not numerically tested the conditions since they are derived for a quadratic PMP cost function. In another more general formulation, Merel et al. (2011) show that a decreasing returns to scale CES function can calibrate exactly to a wider set of elasticities. They also propose that for multiple regions such as in SWAP, the individual region elasticities be allowed to vary as long as the weighted aggregate crop elasticity calibrates to the prior value. This modification will be incorporated in future versions of the SWAP model.

Previous PMP models, such as CVPM, were specified with quadratic PMP land cost functions. Fig. 4 shows a comparison of the exponential PMP cost function and the more frequently used quadratic PMP cost function that implies a linear marginal cost on land. Calibrating a quadratic total cost function subject to a supply elasticity constraint can result in negative marginal costs over a range of low hectares for a specific crop and region. This is inconsistent with basic production theory and can result in numerical difficulties both in the calibration phase and with policy analysis. The exponential cost function is always bounded above zero, by definition, which is consistent with observed costs of production. The marginal factor cost of land has the required first and second order conditions for calibration and minimizes the difference from the prior elasticity value. A second practical advantage is that the exponential cost function often can fit a desired elasticity of supply without forcing the marginal cost of production at low hectares to have unrealistic values. A quadratic PMP cost function, often forces the modeler two choose between an

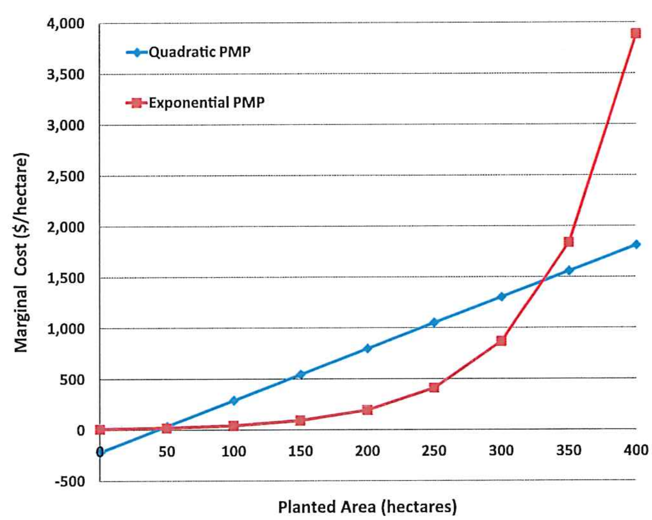

Fig. 4. Comparison of quadratic and exponential PMP land cost functions (adapted from Medellín-Azuara et al., 2010).

unrealistic elasticity, which influences policy response, or an unrealistic initial marginal cost of production. Researchers considering using a quadratic total function should beware of the potential for negative marginal costs.

Formally, in Step IV of PMP calibration we estimate parameters for the exponential cost function. We define the total land cost function as

$$TC(x_{land}) = \delta e^{\gamma x_{land}} \tag{15}$$

where δ and γ are the intercept and the elasticity parameter for the exponential land area response function, respectively. These parameters are from a regression of the calibration shadow values on the observed quantities, restricted by the first order conditions, and elasticity of supply for each crop group from previous studies. For clarity, consider a single PMP cost function within a single region for a specific crop, defined as

$$MC = \frac{\partial TC}{\partial x_{land}} = \delta\gamma e^{\gamma x_{land}}, \tag{16}$$

where marginal cost equals cash cost plus marginal opportunity cost. The acreage supply elasticity, η, is

$$\eta = \frac{\partial x_{land}}{\partial TC}\frac{TC}{x_{land}}, \tag{17}$$

where

$$\frac{\partial x_{land}}{\partial TC} = \frac{1}{\delta\gamma e^{\gamma x_{land}}}. \tag{18}$$

Simplifying and noting that the logarithmic version of the equation is linear,

$$\ln(\eta\delta\gamma x_{land}) + \gamma x_{land} = \ln(R). \tag{19}$$

Thus, two conditions, Equations (16) and (19), must be satisfied at the calibrated (observed) base level of land use. The former is the PMP condition and holds with equality, the latter is the elasticity condition which we fit by least-squares.

The test at this stage of calibration is to calculate the deviation of the marginal PMP cost at the base land allocation from the shadow value of the corresponding calibration constraint, λ^C_{land}, derived in Step II. If deviations are more than a few percentage points in this test, the model does not accurately calibrate, usually due to a non-optimal solution in the least squares fit for the parameters, or an unduly restrictive elasticity constraint on the estimation.

2.2.5. Step V: calibrating demands for endogenous prices

We include endogenous prices through downward sloping demand functions for all crops in SWAP. This represents the consumer side of the market and provides a mechanism for calculation of consumer surplus in the model. As such, the objective function is to maximize the sum of producer and consumer surplus.

We define a subroutine to estimate a statewide demand function for each crop based on the California crop demand elasticity as estimated by Green et al. (2006). We specify the model with linear California-specific crop demand functions. The demand curve represents consumer's willingness-to-pay for a given level of crop production. All else constant, as production of a crop increases, the willingness-to-pay for additional production is expected to fall and to clear the market the price must also fall. The extent of the price decrease depends on the elasticity of demand or, equivalently, the price flexibility. The latter refers to the percentage change in crop price due to a percent change in production given a perfectly competitive market.

We account for regional price differences in the California statewide demand functions. Crop demand includes both in-state and out of state demands for California crops. The statewide demand functions are defined using a base price and regional prices may include deviations from that base price. The state-wide market price of each crop is assumed constant across regions in the state. Regional deviations from the base reflect variations in distance from markets, production contracts, crop quality, variety, harvest season, and other factors.

Production shares by region and price flexibilities of demand are the relevant data needed to calibrate the demand functions. The price flexibilities are based on earlier work for the CVPM model (Reclamation, 1997). We specify a linear inverse-demand function with two parameters, for crop i in region g, defined as

$$p_i = \xi\alpha_i^1 - \alpha_i^2\left(\sum_g\sum_j y_{gij}\right). \tag{20}$$

The crop price is p_i and parameters α_i^1 and α_i^2 represent the intercept and slope of the crop-specific inverse demand curve, respectively. The parameter ξ is a potential parallel shift in demand due to exogenous factors. We calculate the California price for crop i by weighting the regional observed prices v_{gi} by the fraction of region g in the statewide production. Proportion of production (pp_{gi}) is defined as

$$pp_{gi} = \frac{\tilde{y}_{gi}}{\sum_g \tilde{y}_{gi}}, \tag{21}$$

where $\tilde{y}_{gi}$ is the base production. The weighted California price is consequently defined as

$$wp_i = \sum_g v_{gi}pp_{gi}. \tag{22}$$

The regional marketing cost is the difference between the observed regional price (base) and the calculated California crop price. This reflects differences in price which can be attributed to various region-specific differences discussed above and is defined as

$$rmc_{gi} = v_{gi} - wp_{gi}. \tag{23}$$

Given the above definitions, we can calculate the parameters of the inverse demand functions. For a given price flexibilities (χ_i), the slope parameter is

$$\alpha_i^2 = \frac{\chi_i wp_{gi}}{\sum_g \tilde{y}_{gi}}. \tag{24}$$

Consequently, the intercept is

$$\alpha_i^1 = wp_i - \alpha_i^2\sum_g \tilde{y}_{gi}. \tag{25}$$

The test at this stage is to substitute the regional production quantities into Equation (20) and check to see if the equilibrium price adjusted by the regional marketing cost calibrates closely, within a few percentage points, to the regional price.

2.2.6. Step VI: a calibrated non-linear optimization program

The last step in SWAP calibration combines the calibrated functions into a non-linear optimization program. This base program does not include a policy shock and is used to ensure that the calibrated model reproduces observed base year conditions. We include endogenous price determination, agronomic constraints, and resource constraints in the program. With endogenous prices, the objective function is to maximize the sum of producer and consumer surplus.

$$\begin{aligned}\underset{x_{gij},wat_{gw}}{\text{Max}}\; PS+CS = & \sum_i\left(\xi\alpha_i^1\left(\sum_g y_{gi}\right)+\frac{1}{2}\alpha_i^2\left(\sum_g y_{gi}\right)^2\right)\\ & +\sum_g\sum_i\left(rm_{gi}\left(\sum_j y_{gi}\right)\right)\\ & -\sum_g\sum_i\left(\delta_{gi}\exp\left(\gamma_{gi}x_{gi,land}\right)\right)\\ & -\sum_g\sum_i\left(\omega_{gi,supply}x_{gi,supply}+\omega_{gi,labor}x_{gi,labor}\right)\\ & -\sum_g\sum_w\left(\varpi_{gw}wat_{gw}\right). \qquad (26)\end{aligned}$$

The choice variables are inputs (land, labor, water, and other supplies) for each region g and crop i, in addition to total regional water use by source. The first term of the objective function, Equation (26), is the sum of gross revenue plus consumer surplus for all crops and regions, measured relative to the base crop prices. The second term captures the region-specific gross revenue associated from deviations in regional prices from the base prices (these are denoted regional marketing costs). The third term is the region and crop specific PMP land costs. These include both the direct costs of land reported in the base data and the marginal costs inferred from the shadow values on the resource and calibration constraints. The fourth term accounts for labor and other supply costs across all regions and crops. Finally, the fifth term of the objective function is the sum of irrigation water costs by region, crop, and water source. This term is written separately to emphasize that SWAP includes water costs that vary by source.

We define a convex constraint set with resource, agronomic, and other policy constraints. First, the production technology generates the regional crop production y_{gi} as defined in Equation (7). Resource constraints include regional input constraints,

$$\sum_i x_{gij} \le b_{gj} \quad for \quad j \neq \text{water}, \qquad (27)$$

where b_{gi} is total input available by region. Water constraints are incorporated as a restriction on the total water used by region and source,

$$wat_{gw} \le \text{watcons}_{gw} \qquad (28)$$

and total water input use,

$$\sum_i x_{g,i,water} \le \sum_w wat_{gw}. \qquad (29)$$

SWAP allows for movement along the CES production surface, i.e. substitution between inputs. One intensive margin adjustment commonly observed in agriculture is deficit (stress) irrigation. SWAP endogenously determines potential stress irrigation which is dictated by the shape of the respective CES production function. An upper-bound constraint of 15% stress irrigation (relative to the base condition applied water per hectare) is allowed in the model, to prevent the model from reducing applied water rates below the range normally observed. We define the stress irrigation constraint as

$$\frac{x_{gi,water}}{x_{gi,land}} \ge 0.85\, aw_{gi} \qquad (30)$$

Perennial crops are subject to natural retirement or rotation as yields decline in older stands. The average perennial life ($prenlife_i$) is 25 years for almonds and pistachios, other deciduous, and vine crops in SWAP (UCCE, 2011). Subtropical crops have an average life of 30 years. If the time horizon of analysis exceeds 30 years then we expect that farmers have full flexibility to adjust production decisions, including retirement of orchards and vineyards. In the short run we expect farmers devote resources to preserve perennial stands still in prime bearing years. The SWAP model constrains perennial retirement in the short-run (less than the life of the field) to be a proportion of total land use. The proportion is the short-run horizon in years divided by the perennial life. This implicitly assumes that stand age is uniformly distributed and that only older, lower-bearing, fields will be retired. Formally,

$$x_{g,pren,j} \ge \tilde{x}_{g,pren,j}\left(1-\min\left(1,\frac{yr}{prenlife_{pren}}\right)\right), \qquad (31)$$

where $pren \subset i$ and yr is the number of years of the analysis. Marques et al. (2005) demonstrate a two-stage formulation to more explicitly address permanent and annual crops for a range of water availability conditions.

We also include a regional silage constraint for dairy herd feed in the model. The silage constraint forces production to meet the regional feed requirements of the California dairy herd. For example, each cow consumes 20.5 kg of silage per day and corn grain yields are 11.01 tonnes per hectare thus each cow requires about 0.11 silage hectares per year. Multiplying the silage hectares per cow per year by the number of cows in each region yields the minimum silage requirement. The default model assumes a constant herd size into the future, though additional information about future of herd sizes could be used. This constraint can be excluded if the policy being assessed causes relatively small changes in water supply relative to existing regional supplies. Formally,

$$x_{g,corn,land} \ge \tilde{x}_{g,corn,land}, \qquad (32)$$

where $\tilde{x}_{g,corn,land}$ defines the minimum silage constraint for each region.

Maximizing Equation (26) subject to Equations (27)–(29), where production satisfies Equation (7) by choosing the optimal input allocation for each crop and region yields a unique maximum for the SWAP model. The result of the base model run is used to determine if the model calibrates properly. Constraints defined by Equations (30)–(32) are relaxed in the base model in order to check for proper calibration.

There are three fundamental underlying assumptions which we want to emphasize. First, we assume water is interchangeable among crops in the region. Second, a representative regional farmer acts to maximize annual expected profits, equating the marginal revenue of water to its marginal cost. Third, a region selects the crop mix that maximizes profits within that region. This assumes sufficient levels of water storage and internal water distribution capacity and flexibility.

We use the base program to evaluate the fit of the fully calibrated model. The final test for the fully calibrated model compares the percentage difference in input allocation and production output for the model and the base data. The next stage of testing, test 3 in Fig. 1, compares the value marginal product of inputs and their marginal costs for each regional crop input. This test checks that the calibrated model satisfies the necessary conditions for optimization in the CES model (Howitt, 1995b). Before policy scenarios are run, the elasticities of output supply and input demand should be tested