University of California
Santa Cruz

Mathematical and Econometric Modelling of Farm Labor Demand and Financial Risk Management

A project submitted in partial satisfaction
of the requirements for the degree of

MASTER OF SCIENCE

in

APPLIED MATHEMATICS AND STATISTICS

by

Daniel Ladd

June 2013

The project of Daniel Ladd
is approved:

_____
Professor Marc Mangel, Chair

_____
Professor Raquel Prado

_____
Professor Thomas Wu

_____
Dr. Cameron Speir

_____
Dr. Aaron Mamula

Labor Demand as a Function of Project Water Supply in California's San Joaquin Valley

Daniel Ladd[1], Aaron Mamula[2], and Cameron Speir[2]

[1]Center for Stock Assessment Research, University of California, Santa Cruz

[2]NOAA Fisheries, Southwest Fisheries Science Center, Fisheries Ecology Division

June 2013

Abstract

In this paper, we estimate the effect of changes in the irrigation water supply on farm employment. In particular, we describe the economic impacts of providing additional in-stream flow for threatened fish species in California's Sacramento-San Joaquin Delta. With water a crucial, but often scarce, resource in California's Central Valley its allocation involves substantial opportunity costs. Freshwater from the Sacramento-San Joaquin Delta is a component of the agricultural production function, but also serves as a critical habitat for protected fish species, including Delta smelt and Chinook salmon. We analyze the trade-off in water allocation by estimating the relationship between water deliveries to farms and their demand for labor. A significant increase in farm unemployment can be considered an opportunity cost to society of reserving water to maintain fish habitat. Thus, drawing on existing work in agricultural production economics (Shumway 1983, Shumway and Alexander 1984, Moore and Negri 1992), we specify a multioutput production model, which treats water as a fixed, exogenous input. The model takes the form of a system of input demand (including labor demand) and output supply equations that we derive from a general restricted profit function. Letting the profit function be a normalized quadratic, we identify the system of equations and estimate of elasticity of labor demand with respect to water deliveries. Since the agricultural wage may be an endogenous variable in the labor demand function, we also specify a market-level labor demand function that we estimate using instrumental variables. The data include 30 years of county-level employment, water deliveries, crop prices and quantities, relevant production inputs, and weather variables for seven counties in California's San Joaquin Valley.

1 Introduction

The purpose of this paper is to estimate the economic impact of reductions in surface irrigation water deliveries to farms on agricultural labor demand in California's San Joaquin Valley. In particular, we are interested in the effect of policies that allocate water to environmental protection purposes. Towards this aim, we derive and estimate a theoretically justified and empirically tractable agricultural labor demand equation, where labor demand is a function of water supply. We use this model to asses how farmers' labor demand responds to changes in irrigation water deliveries.

In semi-arid regions, such as the western United States, freshwater is a scarce, but immensely valuable resource. Though urban consumers may have the most inelastic demand for water, a large farming sector can often receive substantially more of the freshwater to use for crop irrigation. Water allocated to support environmental aims (e.g. the preservation of protected species and their habitat) also represents an important and growing competing interest for water. Severe droughts complicate policy decisions further by reducing overall available water. In this paper, we consider California's San Joaquin Valley region, where periods of drought and environmentally-focused water pumping restrictions have simultaneously reduced overall available water deliveries. Given the multiple interests for a finite water supply, any water policy decision involves an opportunity cost that should be well understood. Our labor demand model provides a means to evaluate the economic impacts of water supply allocation policies.

Within the San Joaquin Valley region we consider seven counties that receive irrigation water from two government run water supply projects (see Figure 1 for a map of California showing these seven counties). The region, which produced $14 billion worth of crops in 2010, is largely dependent of irrigation water due to the semi-arid climate. The quantity of water delivered by both projects to our study area has varied over time, with deliveries ranging from less than two million acre-feet (MAF) to over six MAF (see Figure 2). During the two most recent droughts in California, the periods from 1987-1992 and 2007-2009, the Sacramento Valley experienced just 57% and 64% of average 1901-2009 runoff (California Department of Water Resources (DWR) 2010). Substantial declines in water deliveries during these two periods can be observed in Figure 2. The temporal variation in water deliveries is itself heterogeneous across counties. Figure 3 depicts the time series of water deliveries and farm employment for each of the seven counties individually.

In California, water is moved from precipitation rich areas in the north of the state to farms and cities in the drier south. Although 75% of California's water supply originates north of Sacramento, the state capital, 75% of water demand lies below the capital city (Carle 2009, p88). The state's developed water is delivered either to urban consumers or to farms, with the farms receiving a roughly 80% share (Carle 2009, p90). Two water projects transport much of the water delivered to the south: the State Water Project (SWP) and the Central Valley Project (CVP). The SWP was created in the 1950s and is operated by the state of California.

It sources water largely from the Feather River watershed in the Sacramento Valley, which it directs down the Feather River towards the Sacramento-San Joaquin River Delta (Carle 2009, p92). The CVP, a Federally operated project, was signed into law by Franklin Roosevelt during the Great Depression and is one of the world's largest water transportation systems. The CVP diverts water from five dammed California rivers. Much of the CVP water comes from dams in the Trinity River and Sacramento River, which are also diverted to the Delta (Carle 2009, p103). Both projects pump water out of the Delta's more southern and inland regions where it is carried mainly through the California Aqueduct to the southern part of the state. Project water is allocated to regional contractors according to SWP and CVP policies. Each year the SWP and CVP decide how to distribute water based on current water conditions (e.g. precipitation, runoff) and habitat requirements for protected fish species. In especially dry years, such as the most recent Californian drought of 2007-2009, both projects substantially reduce deliveries. When populations of endangered fish species drop to dangerously low levels, Delta water pumping is further constrained.

Changes in water deliveries to farms affect farmers' crop mix decisions, which can impact regional economies through multiple paths. Substantial differences in water intensity and value across crops mean variations in water supplies can shift demands for production inputs and have consequences for farm revenue and profits. Sunding et al. (2002) describe a framework to estimate the effects that reduced water deliveries to farms can have on regional economies. In their conceptual model for analyzing water reductions, they identify farm profits, farm revenues, regional product, and farm employment as potentially impacted economic quantities. Sunding et al. argue that reduced water deliveries will, all else equal, shrink farm profits and revenues by restricting farmers' production possibilities. Labor demand for water-intensive crops may also decrease if farmers respond to lower water supply levels by changing their crop mixes. In a connected economy, effects on both farms' income and input demands will have repercussions for the wider regional economy. In this paper we focus our analysis on the effect of water deliveries on farm employment.

The large quantity of water exported annually from the Delta results in degraded habitat conditions for several protected species of fish, particularly Delta smelt, Chinook salmon, and steelhead trout. These exports on average remove approximately 15 percent of natural inflows from the Delta (Lund et al. 2008). The loss of freshwater reduces the quality of water in the Delta (by increasing temperature, salinity, and the concentration of pollutants) and the quantity of rearing habitat for juvenile fish. Pumping operations change the natural flow regimes in the Delta and thus alter migration routes for fish, such that their survival is greatly reduced as they move through the Delta and out to sea. Fish are also killed directly by entrainment at pumping facilities (NMFS 2009).

Existing studies have also estimated the economic impacts to communities from reductions in irrigation water deliveries to farms. Michael et al. (2009) use two approaches to estimate ex post job losses from water

delivery reductions in California's San Joaquin Valley; one uses County Crop Reports and the other is the Statewide Agricultural Production Model (SWAP). The authors estimates agricultural job losses in 2009 to be 4,500 and 4,725 for the former and latter models, respectively. Frisvold and Konyar (2012) use the U.S. Agricultural Resource Model (USARM) to simulate responses to reductions in irrigation water supply in terms of labor demand, acreage planted and farm income for five Southern Mountain (SM) states and California. The USARM, a crop sector-specific non linear mathematical programming model, predicts that for the five SM states a drop in water deliveries of 25% would reduce regional farm labor demand by 3%. Howe and Goemans (2003) compare the impacts of water rights transfers in the South Platte Basin of Colorado with transfers in the Arkansas River basin of Colorado. Using an input/output model, the authors find direct impacts to employment of 1.29 employees per 1,000 acre-feet in the South Platte and 2.02 in Arkansas River. Including indirect impacts, employment losses are estimated at 1.78 employees per 1,000 acre-feet in the South Platte and 2.57 for Arkansas River. Ward et al. (2006) estimate economic impacts in the Rio Grande Basin from increased water scarcity due to droughts and in-stream flow requirements to protect endangered species. Modeling both the hydrologic conditions and crop producers' decisions, the authors simulate lost economic benefits to agriculture ranging from $56 million in the most severe case to less than $1 million for the most benign scenario. Maneta et al. (2009) design a high-resolution hydro-economic model, which uses a constant elasticity of substitution economic model, to simulate impacts from substantially below average rainfall levels, with respect to farm profits, land use, and agricultural employment. Under two different drought scenarios (25% and 50% reductions in precipitation) and considering four representative farms, they project that farms may substantially reduce employment needed for water-intensive crops and slightly increase labor for less-water intensive crops when faced with a water shortfall.

In our study area, water is moved from precipitation rich areas in the north of California to farms and cities in the drier south. In particular, we estimate both a multioutput production model that incorporates farm labor as an input to production and a market-level labor demand model that accounts for the localized nature of agricultural labor markets. The rest of the paper is organized as follows: Section 2 is a description of our methods, including derivation of our multioutput crop production model and description of our labor market model, Section 3 details the data we use, Section 4 presents the empirical specifications we estimate, Section 5 reports our results, and Section 6 offers concluding comments.

2 Methods

In this section we lay out the two modeling frameworks that we use to evaluate the effect of irrigation water supply on agricultural labor demand. In the first, we first derive a system of output supply and input demands, including labor demand functions, using Hotelling's Lemma and the normalized quadratic profit function. We also use an instrumental variables approach to estimate a market-level demand function because the localized

nature of the labor market may mean the agricultural wage is endogenously determined.

2.1 Derivation of output supply and input demand functions

Since farmers receive water largely from two government projects (CVP and SWP) as exogenous deliveries, their decisions regarding water should be responsive to quantity received, as opposed to price paid. Moore and Negri (1992) and Moore et al. (1994) present models of a production process that utilizes an input that is fixed, but can be allocated to multiple outputs. In the following derivation, we closely follow Moore and Negri (1992) by describing farmers as choosing how to allocate production across multiple crop types when irrigation water is a fixed input.

Let $\mathbf{p}$ be a vector of crop prices, $\mathbf{r}$ be a vector of input prices, $\mathbf{w}$ be a vector of crop-specific water allocations, W be total water delivery quantity, and $\mathbf{x}$ be a vector of exogenous input quantities other than water. Also, let the crop-specific profit function for crop j be denoted π_j. Then, total farm profit, $\Pi(\mathbf{p}, \mathbf{r}, W, \mathbf{x})$, for a farmer producing m crops is the solution to a profit maximization problem with farm-level water allocation as a constraint:

$$\Pi(\mathbf{p}, \mathbf{r}, W, \mathbf{x}) = \max_{w_1,\ldots,w_m} \left\{ \sum_{j=1}^{m} \pi_j(p_j, \mathbf{r}, w_j, \mathbf{x}) : \sum_{j=1}^{m} w_j = W \right\}. \tag{1}$$

We can write the profit function as

$$\begin{aligned} \Pi(\mathbf{p}, \mathbf{r}, W, \mathbf{x}) &= \sum_{j=1}^{m} \pi_j(p_j, \mathbf{r}, w_j^*(\mathbf{p}, \mathbf{r}, W, \mathbf{x}), \mathbf{x}) & (2) \\ &= \sum_{j=1}^{m} \pi_j(\mathbf{p}, \mathbf{r}, W, \mathbf{x}), & (3) \end{aligned}$$

where $w_j^*(\mathbf{p}, \mathbf{r}, W, \mathbf{x}), \mathbf{x})$ is the amount of water the farmer uses to produce crop j. Using Hotelling's Lemma we can derive input demand and output supply functions from equation (3). Specifically, Hotelling's Lemma states that the net supply function for a particular good, $y(p)$, is the partial derivative of the profit function, $\pi(p)$, with respect to the price of the good, p:

$$y(p) = \frac{\partial \pi(p)}{\partial p}. \tag{4}$$

Applying Hotelling's Lemma to equation (3), the net supply function for commodity j, y_j is

$$y_j(p_j, \mathbf{r}, W, \mathbf{x}) = \frac{\partial \Pi(\mathbf{p}, \mathbf{r}, W, \mathbf{x})}{\partial p_j}. \tag{5}$$

To estimate the system of output supply and input demand equations in (5), we need a tractable functional form. Following Shumway (1983) and Moore and Negri (1992), among others, we model the profit function using the normalized quadratic, which can be written

$$\Pi(\mathbf{p}, \mathbf{r}, W, \mathbf{x}) = a + \mathbf{B}\tilde{\mathbf{P}} + \frac{1}{2}\tilde{\mathbf{P}}'\mathbf{C}\tilde{\mathbf{P}}, \tag{6}$$

where $\tilde{\mathbf{P}} = [\tilde{\mathbf{p}}, \tilde{\mathbf{r}}, W, \mathbf{x}]$, $\tilde{\mathbf{p}}$ is the vector of crop prices divided by a numeraire crop price , $\tilde{\mathbf{r}}$ is the vector of input prices divided by the same numeraire crop price, and $\{a, \mathbf{B}, \mathbf{C}\}$ are parameters. The normalized quadratic profit function generates output supply and input demand functions of the form:

$$y_j(\mathbf{p}, \mathbf{r}, W, \mathbf{x}) = \alpha^j + \delta^j W + \boldsymbol{\beta}^j \cdot \mathbf{p} + \boldsymbol{\gamma}^j \cdot \mathbf{r} + \boldsymbol{\varepsilon}^j \cdot \mathbf{x}. \tag{7}$$

These equations have the advantage of being linear in parameters and linearly homogeneous.

2.2 Labor demand functions

Using equation (7), and letting r_l be the wage, the labor demand function for the ith farmer, y_l^i is

$$\begin{aligned} y_l^i(\mathbf{p}^i, \mathbf{r}, W^i, \mathbf{x}^i) &= -\frac{\partial \Pi(\mathbf{p}^i, \mathbf{r}, W^i, \mathbf{x}^i)}{\partial r_l} \\ &= \alpha_l + \delta_l W^i + \boldsymbol{\beta}_l \cdot \mathbf{p}^i + \boldsymbol{\gamma}_l \cdot \mathbf{r} + \boldsymbol{\varepsilon}_l \cdot \mathbf{x}^i, \end{aligned} \tag{8}$$

where the negative sign is present because demand is the opposite of net supply.

Although the labor demand function in equation (8) is derived for a single, profit-maximizing farmer, we estimate labor demand at a more aggregated level, such as the county level. By assuming that farmers within a given region are homogeneous, the regional-level labor demand function, $Y_l^r(\mathbf{p_r}, \mathbf{r}, W_r, \mathbf{x_r})$, is of the same form as equation (8):

$$\begin{aligned} Y_l^r &= \sum_i y_l^i(\mathbf{p}^i, \mathbf{r}, W^i, \mathbf{x}^i) \\ &= \sum_i \alpha_l + \delta_l W^i + \boldsymbol{\beta}_l \cdot \mathbf{p}^i + \boldsymbol{\gamma}_l \cdot \mathbf{r} + \boldsymbol{\varepsilon}_l \cdot \mathbf{x}^i \\ &= \alpha_l^r + \delta_l^r W^r + \boldsymbol{\beta}_l^r \cdot \mathbf{p}^r + \boldsymbol{\gamma}_l^r \cdot \mathbf{r} + \boldsymbol{\varepsilon}_l^r \cdot \mathbf{x}^r, \end{aligned} \tag{9}$$

where $\mathbf{p}^r = \bar{\mathbf{p}}_i$, $W^r = \sum_i W_i$, and $\mathbf{x}^r = \bar{\mathbf{x}}_i$. We can estimate this equation with county level data to obtain county labor demand functions and, in particular, identify the response between water deliveries and labor demand, δ_l^r.

2.3 Market-level labor demand model

In the labor demand function in equation (9), for the input price coefficient vector, $\boldsymbol{\gamma}_l^r$, to be consistent, the input prices, $\mathbf{r}$, must be exogenous. However, it is conceivable that farms may not be purely price-takers in the market for agricultural labor. Then, we can represent the labor market as a coupled labor supply function, L^s, and labor demand function, L^d:

$$L^s = \beta_0^s + \beta_1^s Wage_{ag} + \beta_2^s \mathbf{Z^s} + u \tag{10}$$

$$L^d = \beta_0^d + \beta_1^d Wage_{ag} + \beta_2^d \mathbf{Z^d} + v, \tag{11}$$

where $Wage_{ag}$ is the agricultural wage rate, $\mathbf{Z^d}$ is a vector containing W, $\mathbf{p}$, $\mathbf{r}$ (excluding wage price), and $\mathbf{x}$, $\mathbf{Z}^s$ is a vector of exogenous determinants of labor supply, and both u and v are disturbance terms. Imposing the equilibrium condition that labor demand equals labor supply, or $L = L^d = L^s$, we can solve equations (10) and (11) for L and $Wage_{ag}$:

$$L = \beta_0^s + \frac{\beta_1^s}{\beta_1^s - \beta_1^d}(\beta_0^d - \beta_0^s + \beta_2^d \mathbf{Z^d} - \beta_2^s \mathbf{Z^s} + v - u) + \beta_2^s \mathbf{Z^s} + u \tag{12}$$

$$Wage_{ag} = \frac{1}{\beta_1^s - \beta_1^d}(\beta_0^d - \beta_0^s + \beta_2^d \mathbf{Z^d} - \beta_2^s \mathbf{Z^s} + v - u). \tag{13}$$

Since v, the disturbance term from the labor demand equation in (11), appears on the right hand side of equation (13), the agricultural wage is, by definition, endogenous. Therefore, Ordinary Least Squares (OLS) will produce inconsistent estimates for the coefficient $\boldsymbol{\gamma}_l^r$. Intuitively, if v is thought of as a demand shock, in the sense of increased or decreased need for labor, wages may be affected, if the market is local. Recent literature has found evidence of a highly localized labor market in California. Hertz and Zahniser (2013) find rising wages and simultaneously falling employment in California's Kern and Kings counties' agricultural sectors as evidence of a local farm labor shortage. Additionally, Fisher and Knutson (2013) argue that farm labor markets are often localized due a region's specialized labor requirements and worker mobility constraints, such as immigration laws. These studies provide motivation for estimating a market-level labor demand function that treats the agricultural wage as endogeneous.

The problem of endogenous prices in supply and demand models has received much attention in econometrics and one standard approach is to use instrumental variables regression (Baltagi 2007). Using the set of exogenous variables in $\mathbf{Z^s}$ as instruments can alleviate the inconsistency of the wages coefficient. We draw from several recent papers that estimate supply equations for agricultural labor (Buccola et al. 2012, Boucher et al. 2007) to choose our wage instruments. In particular we use the U.S. construction wage, California minimum wage, Mexican GDP and U.S. GDP as instruments, so that the first stage equation is of the form

$$\hat{Wage}_{ag} = \mathbf{\Gamma} \cdot \mathbf{Z} + \epsilon, \tag{14}$$

with $\mathbf{Z} = [\mathbf{Z}^d, \text{U.S. construction wage, California minimum wage, Mexican GDP, U.S. GDP}]$. The second stage equation is

$$L^d = \beta_0^d + \beta_1^d \hat{Wage}_{ag} + \beta_2^d \mathbf{Z^d} + v. \tag{15}$$

3 Data

Our data consist of crop-specific output prices and quantities, input prices, water deliveries, and two weather variables. All data, except for input prices, are taken at the county level and each observation in our data set is a county-year value. Our data cover the 30 years from 1981 to 2010.

3.1 Data sources

Crop specific output quantities are from county agricultural commissioners annual crop reports. Historical records of these reports are maintained by the National Agricultural Statistics Service (NASS) of the United States Department of Agriculture (USDA)[1]. To estimate our multicrop production model, we aggregate crop production in the San Joaquin Valley into eight categories: Nut Orchard, Fruits and Vegetables, Citrus and Fruit Orchard, Field Crops, Cotton, Grapes, Hay, and Pasture. Our aggregation scheme follows similar logic to the crop categories in the agricultural commissioners annual reports and the crop categories used in Statewide Agricultural Production (SWAP) model (Howitt et al. 2001). Our eight-crop categorization is a slightly more aggregated version the the SWAP scheme, which has 12 categories. For example, where they give grain and sugarbeets their own categories, we place both these crops in our Field Crops category. A complete listing of our crop categorization scheme is shown in the Appendix (Tables A1 and A2) using the individual crop names as they appear in the NASS files. Crop category specific output prices are derived from the crop reports by dividing the reported sales revenue by the reported production quantity for each aggregate crop category. In a small number of cases, production quantities are not reported for specific crops in the annual crop reports and we exclude these from price computations and estimation. The proportions of total nominal value for the whole study period that these exclusions represent are given in Appendix Table A3.

For the labor input, we use farm employment data from the BEA Local Area Personal Income (LAPI) Database [2] listed in tables CA25 and CA25N. Specifically, the LAPI database provides the total number of full-time and part-time workers involved in the direct production of crop or livestock commodities for each county in California.

Input prices for seed, fertilizer and fuels are from NASS price indices. They are national indices for prices paid by farmers for the various inputs. The data (listed as: Prices Paid - Fertilizer Index, Prices Paid - Seed Index, Prices Paid - Fuel Index) are available for years 1997-2013 from the NASS Quick Stats database [3]. Since our study time period is 1981-2010, we obtained data from earlier years via a special tabulations request from NASS.

We analyze the agricultural production effects of water deliveries from two projects, the California State Water Project and the federal Central Valley Project. SWP deliveries are reported in California's Department of Water Resources Bulletin 132 [4]. CVP deliveries are reported in the Bureau of Reclamation CVP Ratebook Schedule A-14 [5]. Annual deliveries are reported by irrigation district. To derive deliveries at the county level

[1] Annual crop reports for all California counties can be downloaded at `www.nass.usda.gov/Statistics_by_State/California/Publications/AgComm/Detail/index.asp`

[2] `http://www.bea.gov/iTable/index_regional.cfm`

[3] http://quickstats.nass.usda.gov

[4] `http://www.water.ca.gov/swpao/bulletin.cfm`

[5] `http://www.usbr.gov/mp/cvpwaterrates/ratebooks/irrigation/2013`

we overlay maps of county boundaries on irrigation district boundaries[6]. In cases where irrigation districts span multiple counties, we calculate the percentage of district land area in the two (or more counties) and allocate water deliveries for that district to each county according to the proportion of district land in the county.

We use two variables to control for differing weather conditions across counties and years: cooling degree days and precipitation. Cooling degree days are the number of days with an average temperature above 65 degrees Fahrenheit. Precipitation is the total recorded precipitation (in millimeters). For both variables, the values used in the analysis are the summed totals between April and September (inclusive) as a single annual observation. These data are from the NOAA National Climate Data Center[7]. A list of weather stations used for each county is included in Appendix Table A4.

3.2 Data summary

The seven counties in our study area produced over $14 billion of crops in 2010 (see Figure 4) on 9.4 million acres of land. Table 1 lists harvested acreage, production value and production quantities for the eight crop categories produced in 2010. By production value, Nut Orchard, Citrus and Fruit orchard, and Fruits and Vegetables are the region's most valuable crops and account for 68% of farm revenue in 2010. In Table 2, we present county-specific data on total production value (of all eight crop categories), harvested acreage, production quantity, farm employment, and water deliveries. Fresno, Kern and Tulare counties were the top three counties for all five statistics in 2010.

From 1981 to 2010 revenues have increased from $8 billion (in 2010 dollars) to $14 billion, slightly below the high of over $16 billion in 2007. Figures 5-7 depict the tons produced and corresponding value for Nut Orchard, Citrus and Fruit Orchard, and Fruits and Vegetables, the three most valuable crop groups (in order). Production quantities of Nut Orchard crops, a relatively high-value category, have multiplied about eightfold between 1981-2010, including an 750,000 ton increase during 2007-2008, the largest two-year increase on record. While farmers have increased output of Citrus and Fruit Orchard crops by over one million tons from 1981 to 2010, production and revenues were volatile during the drought periods. The region has also seen its farmers more than triple production quantities of Fruit and Vegetable crops to over 13 million tons.

4 Empirical specification

We estimate the system of input demand and output supply functions in equation (7) for eight outputs (the eight crop categories) and one input (labor). We conduct this estimation at three different spatial resolutions: (1) all seven counties together; (2) Fresno, Kings, and Kern counties together; (3) Fresno, Kings, and Kern

[6]GIS-ready irrigation district boundaries can be downloaded from `http://www.atlas.ca.gov/download.html`

[7]`http://www.ncdc.noaa.gov/land-based-station-data`

counties individually. We single out Fresno, Kings, and Kern counties because they are particularly reliant on delta exports. In the first two spatial scales, we add county dummy variables to control for county-specific fixed effects. We obtain estimates for nine equations in each of the five systems (45 equations estimated total). However, we report just the labor demand equation from all of the systems, with the exception of the Fresno county system, for which we present the entire system[8]. Much of the production literature (Shumway 1983, Shumway and Alexander 1988, Moore and Negri 1992, Moore et al 1994) estimate these systems with Seemingly Unrelated Regression, which yields the same coefficient point estimates as OLS when the cross-equation error terms are uncorrelated, but smaller standard errors. However, we use OLS because we were unable to obtain numerically stable results using SUR.

Given the localized nature of agricultural labor markets (Hertz and Zahniser 2013, Fisher and Knutson 2013), it is likely that the agricultural wage is simultaneously determined by labor demand and supply. We, therefore, also estimate a market-level labor demand model, with the U.S. construction wage, California minimum wage, U.S. GDP and Mexico GDP as instruments. The instrumental variables regression is conducted as a two-stage least squares procedure, with equations (14) and (11) as the first and second stage equations, respectively. We estimate this market-level labor demand function on the same three spatial scales as with the systems of equations approach.

Table 3 summarizes the six empirical specifications described above. In all specifications, we use a one-year lag for crop and input prices to better capture farmers' price expectations. Shumway and Alexander (1988) also use lagged prices, finding that they outperformed in a price forecasting exercise they conducted.

5 Results

Estimation results for labor demand functions from the multicrop production model are shown in Table 4. Of the five deliveries coefficients, those in the seven and three county systems (0.0023 and 0.002, respectively) are statistically significant. The lack of statistical significant for the deliveries coefficients at the county level may reflect a lack of statistical power, since we have 30 observations at the county level, but we pool county observations to get 210 observations for the seven county system. The seven county deliveries coefficient of 0.0023 implies that a 1,000 acre-feet decrease in water deliveries is associated with a loss of 2.3 jobs. To put this number into context, a loss of 456,515 acre-feet of water deliveries (Fresno county's total water received in 2010) would decrease employment by 1,050 jobs, or about 5% of the county's agricultural labor force. Fresno's 2010 water delivery amount is also close to the 500,000 acre-feet not diverted in 2009, specifically for ecological protection purpose (DWR 2010). Thus, we find evidence that decreased water deliveries are associated with a moderate loss of agricultural jobs. Table 5 shows the results for the system of eight crop supply functions for

[8]Results for estimations not reported here are available from the authors upon request

Fresno county.

That the wage coefficient in the seven-county specification is positive and highly statistically significant is likely the result of the endogeneity of the agricultural wage discussed in section 2.3. In Table 6, the results of the instrumental variables estimation appear to mitigate the endogeneity problem, where four of the five estimations show the theoretically correct sign on the wage coefficient, including the seven county estimation, though none are statistically significant. The deliveries coefficients in the seven and three county estimations are 0.0022 and 0.0018, both slightly smaller but statistically significant and consistent with the multicrop production model results. The 95% confidence interval for the market-level model deliveries coefficient in the seven county system is $[0.0011, 0.0033]$. As before the deliveries coefficients for the seven and three county levels are not statistically significant, meaning that we are unable to detect the effect of reduced water deliveries on employment at the county level.

6 Conclusion

Motivated by the potential impacts of reduced water deliveries on agricultural employment in California's San Joaquin Valley, we have estimated a theoretically justified production model and a market-level labor demand model. We find job losses from water deliveries reductions to be around 2 jobs per 1,000 acre-feet for the San Joaquin Valley region. This estimate is consistent with the finding of Sunding et al. (2011) that shrinking deliveries by 1,000 acre-feet results in between 2.2 and 2.4 jobs lost. The authors conclude that 5,000 jobs have been lost due to reductions in water supply to farms by comparing estimated labor demand using 2005 deliveries levels versus those in 2009. We suggest using a five year average (2003-2007) as a baseline with which to compare actual 2009 deliveries. The five year average yields a benchmark of 4.7 MAF, with which to compare 2009 deliveries of 3 MAF. With 1.6 MAF (not equal to the difference of 4.7 and 3 because of rounding) representing the reduction in deliveries, we estimate a loss of 3,287 jobs to the region.

In a Congressional Research Service report, Cody et al. (2009) consider the proportion of water deliveries reduced specifically due to Delta pumping restrictions and conclude that, with a five year average, 20% of the reductions in deliveries are attributable to pumping regulations. Applying the one fifth proportion to the 4,725 jobs and 5,000 jobs estimates of Michael et al. (2009) and Sunding et al. (2011), implies a loss of 945 and 1,000 jobs from environmental restrictions, respectively. Using the same methodology, we estimate job losses at 657 jobs.

Generally, the results indicate that the trade-off between allocating water to farms or to fish is a nuanced one. The agricultural labor market, which serves as the measurement of this trade-off, appears to be complex. Namely, the classical assumption that farms are price-takers in the labor market is likely violated. Recent literature (Hertz and Zahniser 2013, Fisher and Knutson 2013) suggests that the region's agricultural labor

markets are of a local nature and that farm's decisions to increase or decrease hiring has an impact on wages paid. Thus, to appropriately account for the effect of wages on farm labor demand, we explicitly incorporate the endogeneity of the agricultural wage into our model by instrumenting for the wage. Accounting for the endogeneity of the wage is imperative to properly modeling the determinants of wage, which is especially important since we use it to detect the impacts of water policy decisions.

In addition to the complexity of the agricultural labor market, we find that responses to changes in water deliveries are not homogeneous across the seven counties that we study in the San Joaquin Valley. The pooled results for the multicrop production model (Table 4) indicate a loss of 2.3 and 2 jobs associated with a 1,000 acre-foot reduction in water deliveries for the seven and three county systems, respectively. However, estimating the same model for Fresno, Kings and Kern counties individually yields starkly different results for each county. Fresno's deliveries coefficient, at 0.0026, is similar to those at the region-wide levels, but is qualitatively different from the coefficients of -0.0008 and -0.0002 for Kings and Kern counties. From a practical standpoint, we can interpret these results to mean that Kings and Kern counties' labor demand is very unresponsive to changes in water deliveries, all else equal. Thus, the spatial scale is relevant to the magnitude of the impact water deliveries has on labor demand and must be considered when estimating the economic effects of changes in water deliveries

We hope that our results and analysis help inform the debate regarding the allocation of California's Delta water by providing a transparent and theoretically justified model for the relationship between water deliveries and agricultural labor demand.

References

Badi H. Baltagi. *Econometrics*. Springer, 2002.

Stephen R Boucher, Aaron Smith, J Edward Taylor, and Antonio Yúnez-Naude. Impacts of policy reforms on the supply of mexican labor to us farms: new evidence from mexico. *Applied Economic Perspectives and Policy*, 29(1):4–16, 2007.

Steven Buccola, Cheng Li, and Jeffrey Reimer. Minimum wages, immigration control, and agricultural labor supply. *American Journal of Agricultural Economics*, 94(2):464–470, 2012.

David Carle. *Introduction to Water in California: Updated with a New Preface (California Natural History Guides)*. University of California Press, 2009.

Betsy A Cody, Peter Folger, and Cynthia Brougher. California drought: Hydrological and regulatory water supply issues. `http://switchboard.nrdc.org/blogs/dobegi/media/CRS%20CA%20Drought%20Report%20R409791.pdf`, 2009.

Dennis U Fisher and Ronald D Knutson. Uniqueness of agricultural labor markets. *American Journal of Agricultural Economics*, 95(2):463–469, 2013.

George B Frisvold and Kazim Konyar. Less water: How will agriculture in southern mountain states adapt? *Water Resources Research*, 48(5), 2012.

Tom Hertz and Steven Zahniser. Is there a farm labor shortage? *American Journal of Agricultural Economics*, 95(2):476–481, 2013.

Charles W Howe and Christopher Goemans. Water transfers and their impacts: Lessons from three colorado water markets1. *JAWRA Journal of the American Water Resources Association*, 39(5):1055–1065, 2003.

Richard E Howitt, Kristin B Ward, and Siwa M Msangi. Statewide water and agricultural production model: Appendix n, 2001.

Jay Lund, Ellen Hanak, William Fleenor, William Bennet, Richard Howitt, Jeffrey Mount, and Peter Moyle. Comparing futures for the sacramento-san joaquin delta. `http://www.ppic.org/content/pubs/report/r_708ehr.pdf`, 2008.

MP Maneta, M de O Torres, WW Wallender, S Vosti, R Howitt, L Rodrigues, LH Bassoi, and S Panday. A spatially distributed hydroeconomic model to assess the effects of drought on land use, farm profits, and agricultural employment. *Water Resources Research*, 45(11):W11412, 2009.

Jeffrey Michael, Richard Howitt, Josué Medellín-Azuara, and Duncan MacEwan. A retrospective estimate of the economic impacts of reduced water supplies to the san joaquin valley in 2009, 2010.

Michael R Moore and Ariel Dinar. Water and land as quantity-rationed inputs in california agriculture: Empirical tests and water policy implications. *Land Economics*, pages 445–461, 1995.

Michael R Moore, Noel R Gollehon, and Marc B Carey. Multicrop production decisions in western irrigated agriculture: the role of water price. *American Journal of Agricultural Economics*, 76(4):859–874, 1994.

Michael R Moore and Donald H Negri. A multicrop production model of irrigated agriculture, applied to water allocation policy of the bureau of reclamation. *Journal of Agricultural and Resource Economics*, pages 29–43, 1992.

California Department of Water Resources. California drought: An update. `http://www.water.ca.gov/waterconditions/drought/docs/DroughtReport2010.pdf`.

National Marine Fisheries Service. Biological opinion and conference opinion on the long-term operations of the central valley project and state water project. `http://www.swr.noaa.gov/ocap/NMFS_Biological_and_Conference_Opinion_on_the_Long-Term_Operations_of_the_CVP_and_SWP.pdf`.

C Richard Shumway. Supply, demand, and technology in a multiproduct industry: Texas field crops. *American Journal of Agricultural Economics*, 65(4):748–760, 1983.

C Richard Shumway and William P Alexander. Agricultural product supplies and input demands: regional comparisons. *American Journal of Agricultural Economics*, 70(1):153–161, 1988.

David Sunding, David Zilberman, Richard Howitt, Ariel Dinar, and Neal MacDougall. Measuring the costs of reallocating water from agriculture: A multi-model approach. *Natural Resource Modeling*, 15(2):201–225, 2002.

L David Sunding, Kate Foreman, and Maximilian Auffhammer. Water and jobs: The role of irrigation water deliveries on agricultural employment, 2011.

Frank A Ward, Brian H Hurd, Tarik Rahmani, and Noel Gollehon. Economic impacts of federal policy responses to drought in the rio grande basin. *Water Resources Research*, 42(3):W03420, 2006.

7 Tables

Table 1: Crop mix summary statistics for the San Joaquin Valley of California in 2010

Crop Group	Production value ($)	Harvested acres	Production (tons)
Nut orchard	3,963,826,000	880,983	2,480,330
Citrus and fruit orchard	3,172,047,000	408,739	4,170,555
Fruits and vegetables	2,395,462,000	443,949	13,521,280
Grapes	2,346,625,000	421,181	3,470,766
Field	941,744,000	1,312,422	22,287,642
Cotton	571,638,000	292,200	193,230
Hay	524,489,000	683,026	4,377,260
Pasture	107,239,000	4,915,320	4,915,320
Total	14,023,070,000	9,357,820	55,416,383

Table 2: Crop production summary statistics for individual counties of San Joaquin Valley in 2010

County	Production value ($)	Harvested acres	Production (tons)	Ag. jobs	Water deliveries (acre-feet)
Fresno	4,443,618,000	1,869,969	13,145,980	20,500	456,515
Kern	3,124,122,000	2,270,004	7,504,270	16,986	1,850,316
Tulare	2,542,637,000	1,627,658	12,211,120	16,660	751,385
Stanislaus	1,142,977,000	1,040,811	7,047,890	9,792	41,839
Merced	1,044,363,000	1,144,660	7,032,863	8,469	45,724
Madera	867,264,000	647,690	2,432,170	5,358	338,567
Kings	858,089,000	757,028	6,042,090	4,300	160,401
Total	14,023,070,000	9,357,820	55,416,383	82,065	3,644,747

Table 3: Empirical Specifications

Multicrop Production Model (with prices assumed exogenous)		
Counties	Pooled?	Instrumental Variables?
All seven	yes	no
Fresno, Kings, and Kern	yes	no
Fresno, Kings, and Kern	no	no
Market-level Labor Demand Model (with wage instruments)		
Counties	Pooled?	Instrumental Variables?
All seven	yes	yes
Fresno, Kings, and Kern	yes	yes
Fresno, Kings, and Kern	no	yes

Table 4: Farm labor demand from the multicrop production model

	7 counties	3 counties	Fresno county	Kings county	Kern county
Variable			Coefficients		
Deliveries	0.0023****	0.0020***	0.0026	-0.0008	-0.0002
Nut Orchard Price	-194***	-315**	-528***	-32	-201
Fruits and Veg. Price	-172	-2202**	-5538	222	-4058
Citrus and Fruit Orchard Price	-455*	-182	-207	-166	146
Field Price	2088	-1353	-3172	-1202	-9564
Cotton Price	-46	46	566	-18	1000
Grapes Price	616**	1068*	1513	105	-264
Hay Price	2630*	7390**	7764	1544	6572
Pasture Price	1316	35991**	160152	4045	19526
Fertilizer Price	2809	1767	9520	184	2372
Fuel Price	-5418****	-4561*	-11627**	-2450**	-9064*
Wage	138114****	-5715	-27478	17564	8075
Cooling Degree Days	-0.16	-0.55***	-0.44**	0.02	0.16
Precipitation	0.03	-0.02	3.57**	0.08	0.69
Year	62	-48	186	-19	343

*p<0.10; **p<0.05; ***p<0.01; ****p<0.001

Table 5: Crop Supply Functions from the multicrop production model (Fresno County only)

	Nut Orch.	F. & V.	C/F Orch.	Field	Cotton	Grapes	Hay	Past.
Variable	Coefficients							
Deliveries	-0.0471	-0.4108	0.1211	-0.8355	0.0165	-0.0852	0.0089	0.0425*
Nut Orch.	5897	16652	7795	99308	4290	-23181	5141	357
F. & V.	184766	1746209	100722	6863274	-26493	-459846	411202	161406**
C/F Orch.	-35353	-85646	22288	-769629	4186	3178	-35291	-9707
Field	-72434	1029825	165918	4180235	37714	307732	-64962	39564
Cotton	-9552	24541	-26850	-219304	4879	-65005	-47574*	-14059**
Grapes	44882*	-464194*	-47909	-494285	-28089	89750	-40776	3257
Hay	-130395	749465	-86729	-1199689	-48255	1158701	219923	-71801
Past.	5280136	-26679754	-3343636	180066910	-4270699	-2479841	643597	2073013
Fert.	62283	-602535	-447738	-241491	-113197	-27309	226555	123877
Fuel	-135954	-135510	147996	-1724645	38921	-1031967	92071	-75791
Wage	-14126272**	16447341	2567618	-194092415	8829521**	32046236	-484145	-884089
C. D. D.	4.89	-67.16	24.45	-140.51	2.55	-3.68	1.03	4.00
Precip.	-19	-131	35	1300	-54	-479	-65	11
Year	21253*	158948	12033	566669	-14079	-14167	-5176	102

*p<0.10; **p<0.05; ***p<0.01; ****p<0.001

Table 6: Market Level Agricultural Labor Demand: Instrumental Variables Regression Results

	7 counties	3 counties	Fresno county	Kings county	Kern county
Variable			Coefficients		
Deliveries	0.0022****	0.0018**	-0.0083	0.0009	-0.0006
Nut Orchard	-0.90**	-1.45**	2.83**	0.29	0.59
Fruits and Veg.	3.11	-11.44**	88.77	-1.38	16.62
Citrus and Fruit Orchard	-1.60	-0.74	-11.14	0.34	-10.06
Field	2.67	-20.13	8.92	6.42	105.14
Cotton	-0.20	-0.06	5.09	-0.13	-10.25
Grapes	5.83****	10.50****	-4.87	0.32	-0.75
Hay	30.47*	60.43**	-163.34	-2.00	-102.84
Pasture	12.24	161.63	3184.70	0.61	-308.33
Fertilizer Price	-5.14	-9.50	96.31	-5.83	45.73
Fuel Price	-26.59****	-27.79**	-27.16	11.40	42.85
Seed Price	15.40	31.15	-88.66	8.61	-102.29
Wage	-248.32	-655.30	-5392.10	-118.36	3749.05
Cooling Degree Days	-0.09	-0.46**	0.12	-0.01	-0.78
Precipitation	0.03	-0.05	2.62	-0.38	-0.92
Year	48.68	-161.91	1712.15	-41.54	-541.21

*p<0.10; **p<0.05; ***p<0.01; ****p<0.001

8 Figures

Figure 1: Map of California including the seven counties in the study area

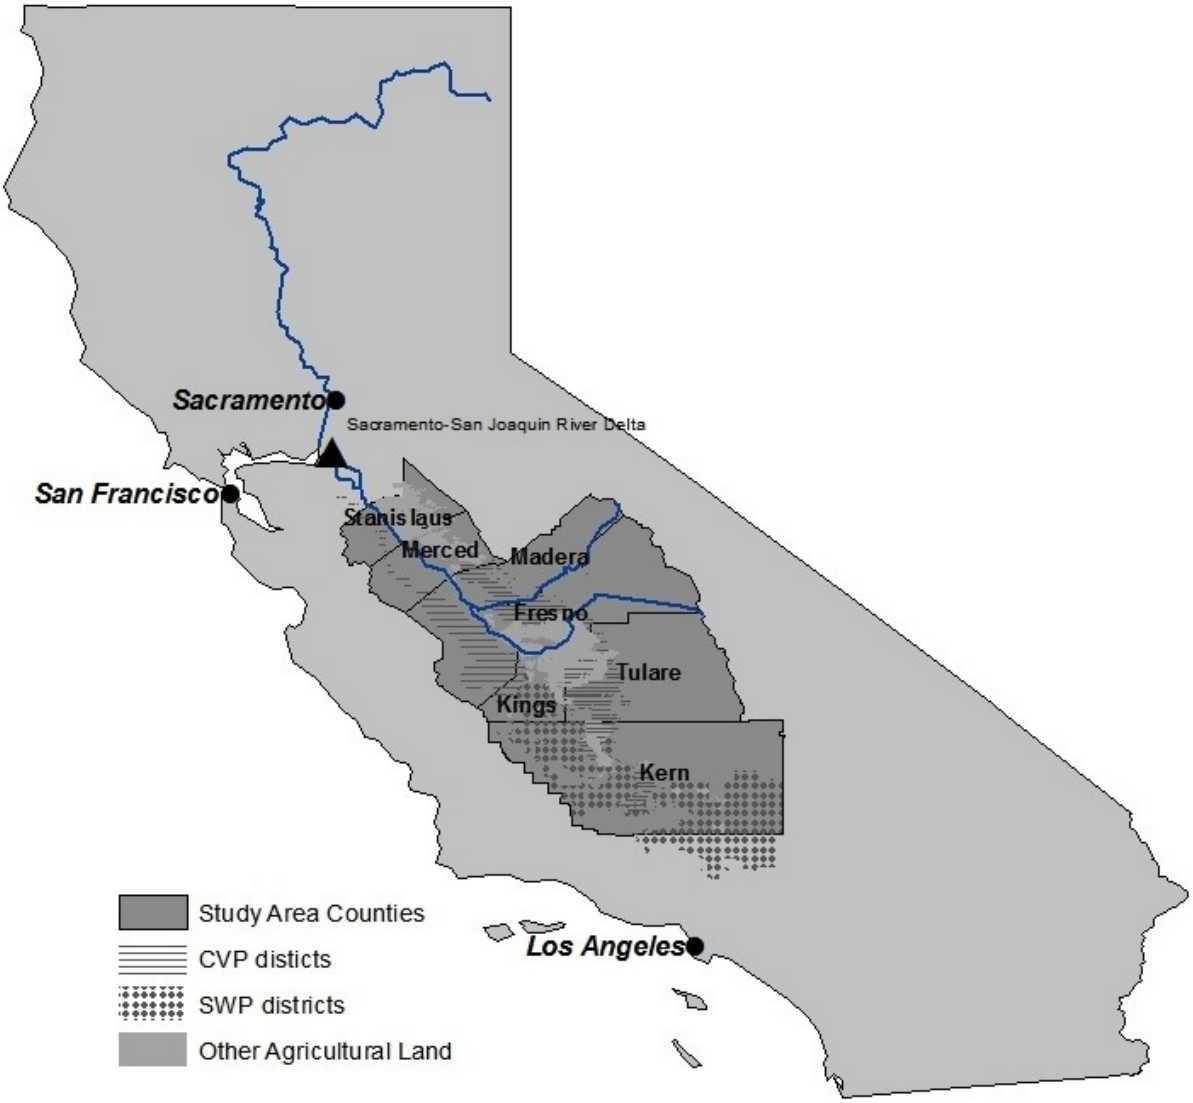

Figure 2: Total project (CVP and SWP) water deliveries and total agricultural employment in the seven-county study, 1981-2010

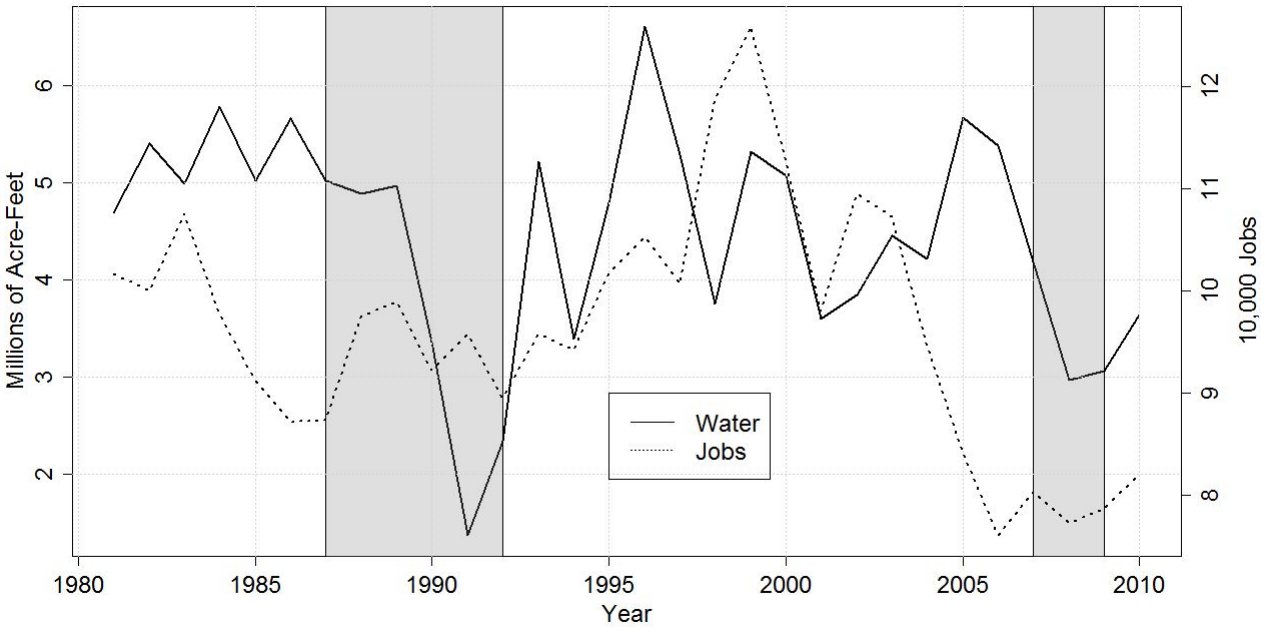

SWP Source: Table B-5B of Appendix B from Bulletin 132-12 available `http://www.water.ca.gov/swpao/docs/bulletin/12/Appendix_B.pdf`. CVP Source: Schedule A-14 of CVP Ratebooks-Irrigation available `http://www.usbr.gov/mp/cvpwaterrates/ratebooks/irrigation/2012/2012_irr_sch_a-14.pdf`. Shaded regions represent droughts.

Figure 3: Project (CVP and SWP) water deliveries and total agricultural employment for each of the seven counties, 1981-2010

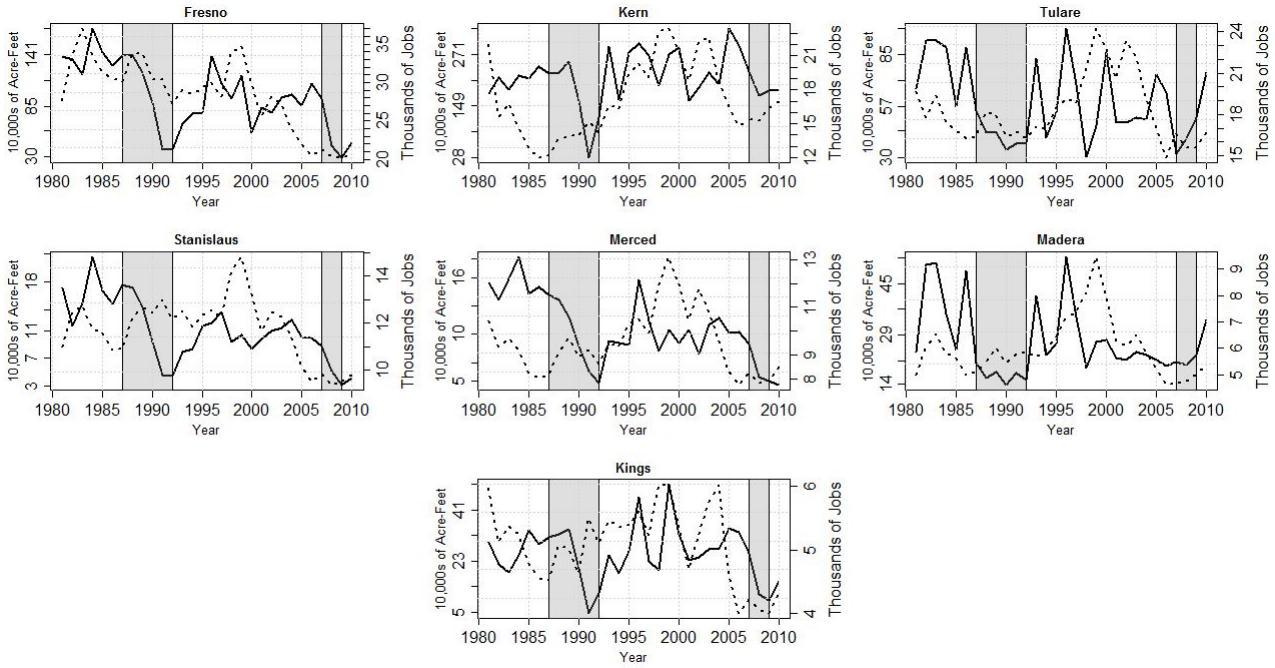

SWP Source: Table B-5B of Appendix B from Bulletin 132-12 available `http://www.water.ca.gov/swpao/docs/bulletin/12/Appendix_B.pdf`. CVP Source: Schedule A-14 of CVP Ratebooks-Irrigation available `http://www.usbr.gov/mp/cvpwaterrates/ratebooks/irrigation/2012/2012_irr_sch_a-14.pdf`. Shaded regions represent droughts.

Figure 4: Total value of agricultural crop production in the San Joaquin Valley of California: 1981-2010 (inflation-adjusted 2010 dollars)

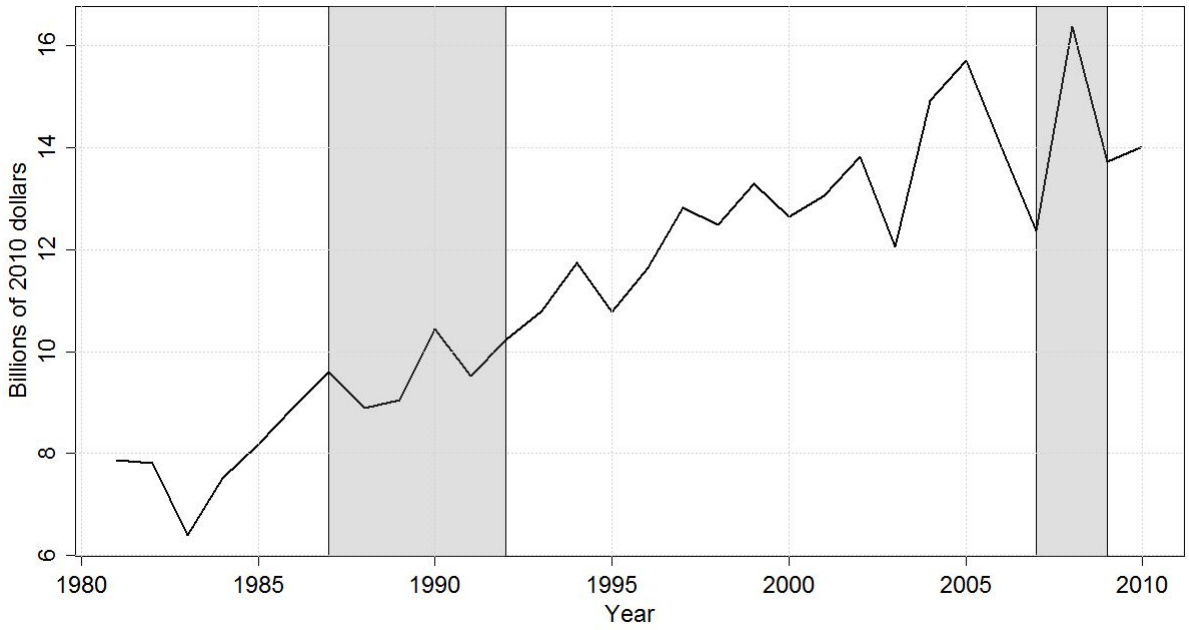

Source: County agricultural commissioners crop reports 1981-2010, available from `www.nass.usda.gov/Statistics_by_State/California/Publications/AgComm/Detail/index.asp`. Shaded regions represent droughts.

Figure 5: Total value and quantity of nut orchard crop production in the San Joaquin Valley of California: 1981-2010 (inflation-adjusted 2010 dollars)

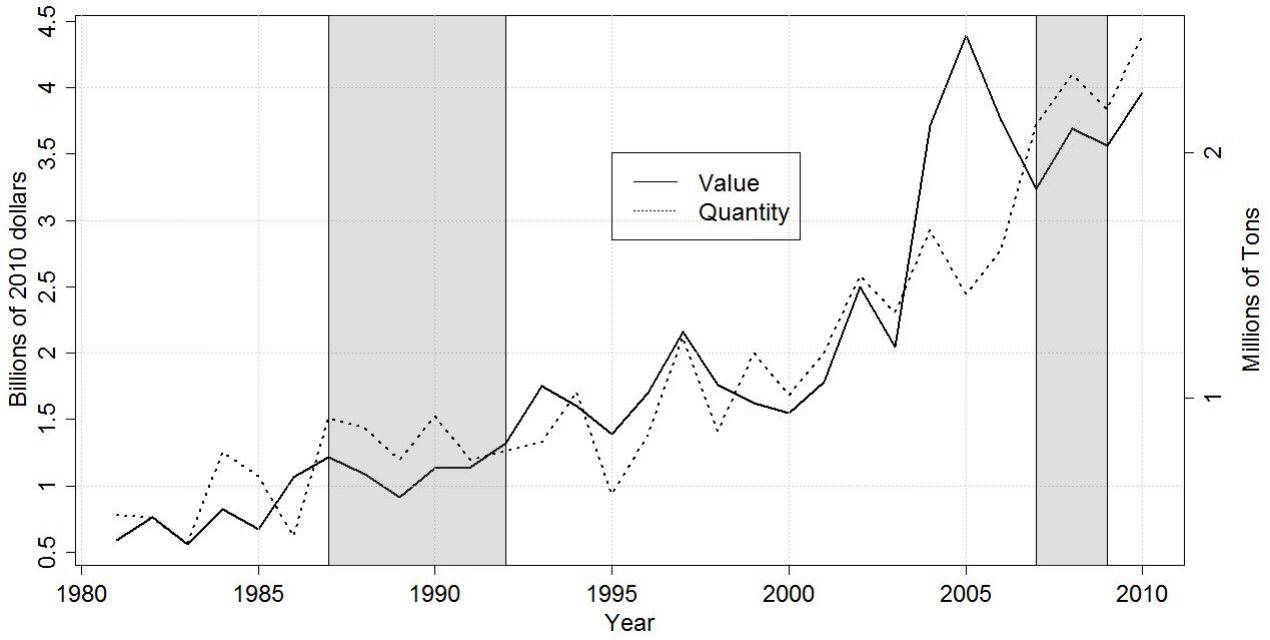

Source: County agricultural commissioners crop reports 1981-2010, available from `www.nass.usda.gov/Statistics_by_State/California/Publications/AgComm/Detail/index.asp`. Shaded regions represent droughts.

Figure 6: Total value and quantity of citrus and fruit orchard crop production in the San Joaquin Valley of California: 1981-2010 (inflation-adjusted 2010 dollars)

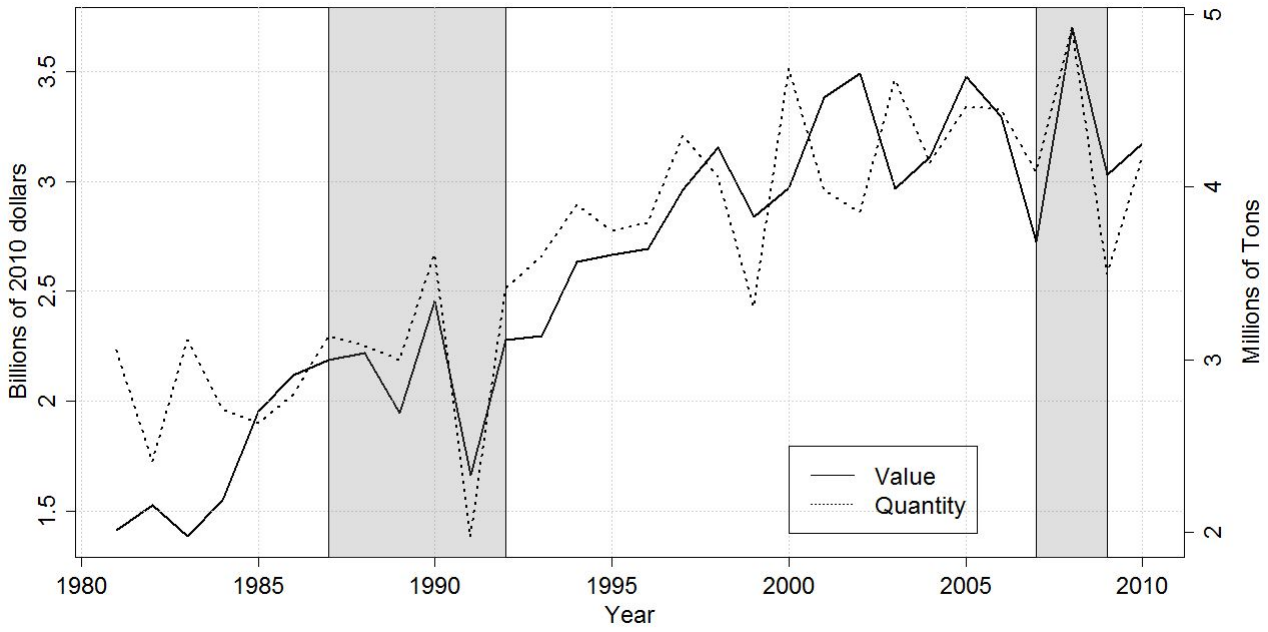

Source: County agricultural commissioners crop reports 1981-2010, available from `www.nass.usda.gov/Statistics_by_State/California/Publications/AgComm/Detail/index.asp`. Shaded regions represent droughts.

Figure 7: Total value and quantity of fruit and vegetable crop production in the San Joaquin Valley of California: 1981-2010 (inflation-adjusted 2010 dollars)

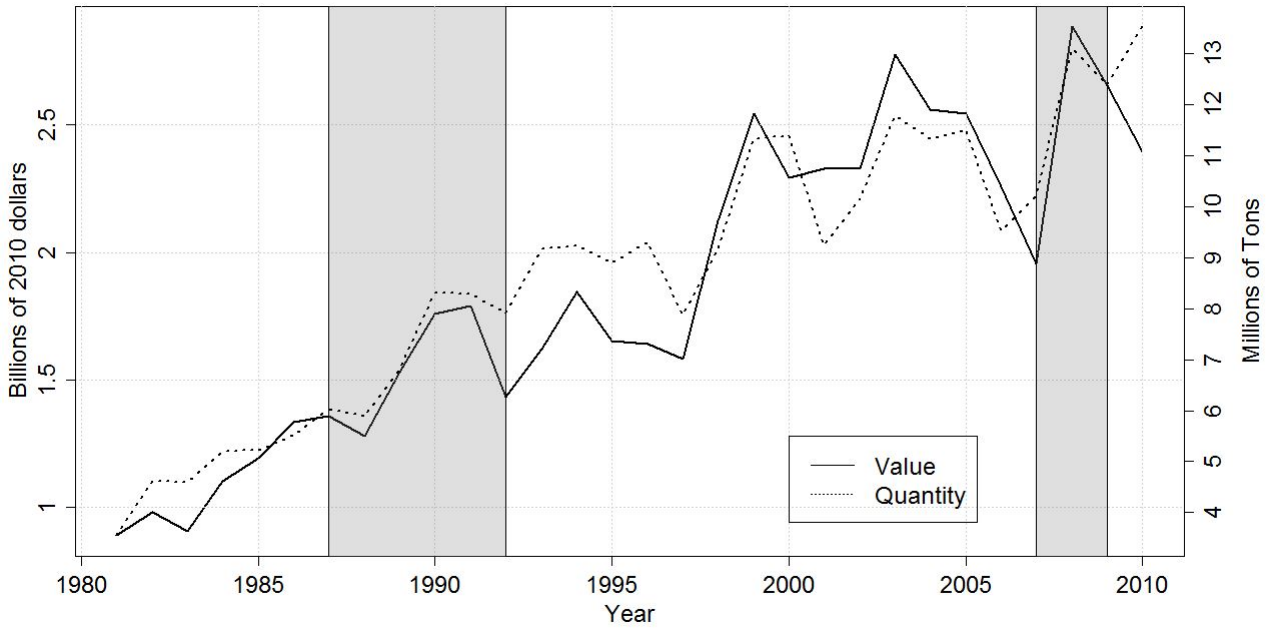

Source: County agricultural commissioners crop reports 1981-2010, available from `www.nass.usda.gov/Statistics_by_State/California/Publications/AgComm/Detail/index.asp`. Shaded regions represent droughts.

9 Appendix

Table A1: Group Grouping Part A

Crop Group	Crop Name
Nut Orchard	ALMONDHULLS,ALMONDSALL,CHESTNUTS,FILBERTS,MACADAMIANUTS,PECANS,PISTACHIOS,WALNUTSBLACK, WALNUTSENGLISH
Fruits and Vegetables	ANISE(FENNEL),ARTICHOKES,ASPARAGUSFRESHMARKET,ASPARAGUSPROCESSING,ASPARAGUSUNSPECIFIED, BEANSFAVA,BEANSFRESHUNSPECIFIED,BEANSSNAPFRESHMARKET,BEANSSNAPPROCESSING, BEANSSNAPUNSPECIFIED,BEETSGARDEN,BROCCOLIFOODSERVICE,BROCCOLIFRESHMARKET, BROCCOLIPROCESSING,BROCCOLIUNSPECIFIED,BRUSSELSSPROUTS,CABBAGECH.&SPECIALTY, CABBAGECHINESE&SPECIALTY,CABBAGEHEAD,CABBAGERED,CARDOON,CARROTSFOODSERVICE, CARROTSFRESHMARKET,CARROTSPROCESSING,CARROTSUNSPECIFIED,CAULIFLOWERFOODSERVICE, CAULIFLOWERFRESHMARKET,CAULIFLOWERPROCESSING,CAULIFLOWERUNSPECIFIED,CELERYFOODSERVICE, CELERYFRESHMARKET,CELERYPROCESSING,CELERYUNSPECIFIED,CHAYOTES,CHIVES,CILANTRO, COLLARDGREENS,CORNCRAZY,CORNPOPCORN,CORNSWEETALL,CUCUMBERS,CUCUMBERSGREENHOUSE, EGGPLANTALL,ENDIVEALL,ESCAROLEALL,GARLICALL,GREENSTURNIP&MUSTARD,HOPS,HORSERADISH,KALE, KOHLRABI,LEEKS,LETTUCEBULKSALADPRODS.,LETTUCEBULKSALADPRODUCTS,LETTUCEHEAD,LETTUCELEAF, LETTUCEROMAINE,LETTUCEUNSPECIFIED,MELONSCANTALOUPE,MELONSCASABA,MELONSCRENSHAW, MELONSHONEYDEW,MELONSUNSPECIFIED,MELONSWATERMELON,MINT,MUSHROOMS,MUSTARD,OKRA,ONIONS, ONIONSGREEN&SHALLOT,PARSLEY,PARSNIPS,PEANUTSALL,PEASEDIBLEPOD(SNOW),PEASGREENFRESHMARKET, PEASGREENPROCESSING,PEASGREENUNSPECIFIED,PEPPERSBELL,PEPPERSCHILIHOT,PIMENTOS,POTATOESALL, POTATOESIRISHALL,POTATOESSWEET,PUMPKINS,RADICCHIO,RADISHES,RAPPINI,RHUBARB,RUTABAGAS, SALADGREENSMISC.,SALADGREENSNEC.,SESAME,SPICESANDHERBS,SPINACHFOODSERVICE, SPINACHFRESHMARKET,SPINACHPROCESSING,SPINACHUNSPECIFIED,SPROUTSALFALFA&BEAN,SQUASH, SWISSCHARD,TAROROOT,TOMATILLO,TOMATOESCHERRY,TOMATOESFRESHMARKET,TOMATOESGREENHOUSE, TOMATOESPROCESSING,TOMATOESUNSPECIFIED,TURNIPSALL,VEGETABLESBABY,VEGETABLESGREENHOUSE, VEGETABLESORIENTALALL,VEGETABLESUNSPECIFIED,WATERCRESS

Table A2: Group Grouping Part B

Crop Group	Crop Name
Citrus and Fruit Orchard	APPLESALL,APRICOTSALL,AVOCADOSALL,BERRIESBLACKBERRIES,BERRIESBLUEBERRIES, BERRIESBOYSENBERRIES,BERRIESBUSHBERRIESUNSPEC.,BERRIESBUSHBERRIESUNSPECIFIED, BERRIESLOGANBERRIES,BERRIESOLALLIEBERRIES,BERRIESRASPBERRIES,BERRIESSTRAWBERRIESFMKT, BERRIESSTRAWBERRIESFRESHMARKET,BERRIESSTRAWBERRIESPROC.,BERRIESSTRAWBERRIESPROCESSING, BERRIESSTRAWBERRIESUNSPEC,BERRIESSTRAWBERRIESUNSPECIFIED,CHERIMOYAS,CHERRIESSWEET, CITRUSBY-PRODUCTSMISC.,CITRUSUNSPECIFIED,DATES,FEIJOA,FIGSDRIED,FRUITS&NUTSUNSPECIFIED, GRAPEFRUITALL,GUAVAS,JOJOBA,KIWIFRUIT,KUMQUATS,LEMONSALL,LIMESALL,NECTARINES,OLIVES, ORANGESNAVEL,ORANGESUNSPECIFIED,ORANGESVALENCIA,PEACHESCLINGSTONE,PEACHESFREESTONE, PEACHESUNSPECIFIED,PEARSASIAN,PEARSBARTLETT,PEARSPRICKLY,PEARSUNSPECIFIED,PERSIMMONS, PLUMCOTS,PLUMS,PLUMSDRIED,POMEGRANATES,POMELO,QUINCE,TANGELOS,TANGERINES&MANDARINS
Field	BARLEYFEED,BARLEYMALTING,BARLEYUNSPECIFIED,BEANSBLACKEYE(PEAS),BEANSDRYEDIBLEUNSPEC., BEANSDRYEDIBLEUNSPECIFIED,BEANSGARBANZO,BEANSKIDNEYRED,BEANSLIMABABYDRY, BEANSLIMAGREEN,BEANSLIMALARGEDRY,BEANSLIMALG.DRY,BEANSLIMAUNSPECIFIED,BEANSPINK, BEANSPINTO,BEANSREDSMALL,BEANSWHITESMALL,BEANSWHITESMALLFLAT,CORNGRAIN,CORNSILAGE, CORNWHITE,COTTONSEEDPLANTING,FIELDCROPBY-PRODUCTS,FIELDCROPSUNSPECIFIED,FOODGRAINSMISC., GUAR,OATSGRAIN,PEASCOWPEA&BLACKEYE,PEASDRYEDIBLE,RICEMILLING,RICESWEET,RICEWILD,RYEGRAIN, SAFFLOWER,SILAGE,SORGHUMGRAIN,SORGHUMSILAGE,SOYBEANS,SUGARBEETS,TRITICALE,WHEATALL
Cotton	COTTONLINTPIMA,COTTONLINTUNSPECIFIED,COTTONLINTUPLAND
Grapes	GRAPESRAISIN,GRAPESTABLE,GRAPESUNSPECIFIED,GRAPESWINE
Hay	HAYALFALFA,HAYGRAIN,HAYGREENCHOP,HAYOTHERUNSPECIFIED,HAYSUDAN,HAYWILD,STRAW
Pasture	PASTUREFORAGEMISC.,PASTUREIRRIGATED,PASTURERANGE

Table A3: Nominal Value Discarded Due to Missing Production Value

Crop	% of Total Nominal Value
Nut Orchard	0.0015%
Fruits and Veg.	1.1773%
Citrus and Fruit Orchard	0.5150%
Field	0.2181%
Cotton	0.0000%
Grapes	0.0000%
Hay	0.0001%
Pasture	0.0090%

Table A4: Weather Stations

County	Weather Station
Glenn	WILLOWS 6 W CA US
Colusa	COLUSA 2 SSW CA US
Sutter	COLUSA 2 SSW CA US
Yolo	WOODLAND 1 WNW CA US
Sacramento	SACRAMENTO 5 ESE CA US
San Joaquin	TRACY CARBONA CA US
Stanislaus	MODESTO CITY CO AIRPORT CA US
Merced	NEWMAN CA US
Madera	MADERA CA US
Fresno	FIVE POINTS 5 SSW CA US
Tulare	VISALIA CA US
Kings	CORCORAN IRRIG DIST CA US
Kern	WASCO CA US

Implied Volatility Value at Risk: a Forward-looking Risk Measure

Daniel Ladd

June 2013

Contents

1 Introduction

In the field of quantitative financial risk management, we often focus on the probability distribution of the value of an asset or liability at some point in the future. Given a future time period of interest, we face two substantial questions: (1) how do we construct such a probability distribution, and (2) how do we use it to risk manage? Though there are multiple answers to the former question, to a great extent the use of Value at Risk (VaR) has been the dominant answer to the latter (Glasserman 2003, p 483). VaR is used to determine capital requirements, but also to set trading risk limits and compensation levels for traders.

1.1 Defining Value at Risk

Given a confidence level, α, and the cumulative distribution function of future returns, $F(r)$, VaR is defined as

$$\text{VaR}_\alpha = \{x : F(x) = 1 - \alpha\}. \tag{1}$$

VaR is simply a downside quantile, typically with $\alpha = 1\%$ or 5%. Also, we call the occurrence of a negative return larger in absolute value that VaR a "VaR event". For example, if today's 95% VaR is \$1,000,000, we expect losses worse than this to happen only 5% of days, and we call any such losses "shortfalls". VaR "waiting times" are the lengths of times between successive VaR events.

Despite its widespread use in banking activities, both academics and practitioners often deem it inadequate and to provide a false sense of security (Glasserman 2003, p 483). However, these criticisms are directed at VaR calculated using historical data, whereas little, if any, research has been conducted on calculations of VaR using alternative data sources. Using historical data to predict future market dynamics (namely volatility) requires the assumption that either we have a clear understanding about the data generating process for market moves or that patterns are apparent and predictable. Market crashes, such as that associated with the most recent financial crisis, are evidence that financial markets exhibit extremely complex and unpredictable behavior. The problem with using historical data for risk management is its inherently backward-looking nature, which forces us to specify how the past relates to the future.

As a forward-looking alternative to using historical data, I propose that we calculate VaR using option-implied volatility. Intuitively, option-implied volatilities are market participants' forecast of

future volatility. New information, be it a just-released gross domestic product report or a proprietary research report, is reflected in implied volatility as soon as investors respond by making option trades. Additionally, any predictive power about future market movements imbedded in historical market data should be present in implied volatility, assuming a rational market. Thus, implied volatility may provide a substantial advantage over historical data because of its continuously updating and forward-looking nature.

1.2 Option-Implied Volatility

Since I use option-implied volatilities to compute what I will call Implied Volatility VaR, I give a brief introduction to it here. In the foreign exchange market, at time t_0, we observe the price that investors are willing to pay for the option, but not the obligation, to exchange euros for dollars (in the case of EURUSD) for the current exchange rate to another party at a future date (i.e. an at-the-money option). At expiration, this contract will have a positive payoff if the dollar has appreciated against the euro since time t_0 and zero payoff otherwise.

The Black-Scholes model is a textbook approach to price options and is used to determine option-implied volatilities. Given the asset's volatility, σ (from a Geometric Brownian Motion process), the interest rate differential between currencies, ρ, and the time to expiration τ, according to Black-Scholes, the value of an at-the-money put option, P, is

$$P(S(t), \tau, \rho, \sigma) = \Phi(-d_2)S(t_0)e^{-\rho\tau} - \Phi(-d_1)S(t), \tag{2}$$

where $S(t)$ is the asset price at time t, Φ is the normal cumulative distribution function, and

$$d_1 = \frac{(\rho + \sigma^2/2)\tau}{\sigma\sqrt{\tau}}, \tag{3}$$

$$d_2 = d_1 - \sigma\sqrt{\tau}. \tag{4}$$

In general, all of the inputs to this model are easy to determine, except one: the asset's volatility. Given that we observe market prices, we can use numerical methods to determine implied volatility, $\tilde{\sigma}$, or the value of σ that will yield the market price using the Black-Scholes formula (Wilmott 1998, p 287). Hence, the name implied volatility. This estimate of volatility, $\tilde{\sigma}$ incorporates all market expectations about future volatility and can be used in place of other estimates of volatility.

2 Portfolio Construction

For the remainder of this paper, the investment portfolio will be comprised of a long position in GBPUSD (i.e. the number of Dollars per Pound) and a short position in EURUSD (i.e. the number of Dollars per Euro). Specifically, the investment entails initially allocating equal amounts, say $1 million, towards both trade positions. Since they are of equal initial value, the two trades make the portfolio USD neutral at construction. However, as the GBPUSD and EURUSD rates drift, so, too, will the relative weights of each currency pair in the portfolio. When the weights are not equal, the portfolio is no longer USD hedged and is exposed to the fluctuations in the value of the dollar.

Holding a long and short position in two positively correlated assets is common of broker/dealers who may wish to limit market risk, while fulfilling multiple client orders. The challenge for this representative broker/dealer is to characterize risk for a portfolio that is generally stable, but may become unexpectedly volatile in extreme market conditions, such as a financial crisis.

In the first time period, the return observed on the portfolio will simply be the day's return on GBPUSD less the return on EURUSD, or the negative of the EURGBP return. However, as the exchange rates change the effective portfolio weights will also vary. Thus, letting r_p be the portfolio return, r_{GBP} be the GBPUSD return, and r_{EUR} be the EURUSD return, the portfolio returns are calculated as

$$r_p(t) = 2\left[\frac{\text{GBPUSD(t-1)}}{\text{GBPUSD(1)}} \Big/ \left(\frac{\text{GBPUSD(t-1)}}{\text{GBPUSD(1)}} + \frac{\text{EURUSD(t-1)}}{\text{EURUSD(1)}}\right) r_{GBP}\right. \tag{5}$$

$$\left. - \frac{\text{EURUSD(t-1)}}{\text{EURUSD(1)}} \Big/ \left(\frac{\text{GBPUSD(t-1)}}{\text{GBPUSD(1)}} + \frac{\text{EURUSD(t-1)}}{\text{EURUSD(1)}}\right) r_{EUR}\right], \tag{6}$$

where the factor of two makes the returns in the beginning time periods comparable to EURGBP returns, since the returns on EURGBP are r_{EUR} less r_{GBP}. Figure 1 and Figure 2 show the relative value of the portfolio and its daily returns for the time period of 3/2006 to 4/2013.

3 Calculating Value at Risk

Although the definition of VaR is simple and precise, how exactly one should calculate it is subject to interpretation and debate. The three VaR models detailed below are similar in the sense that

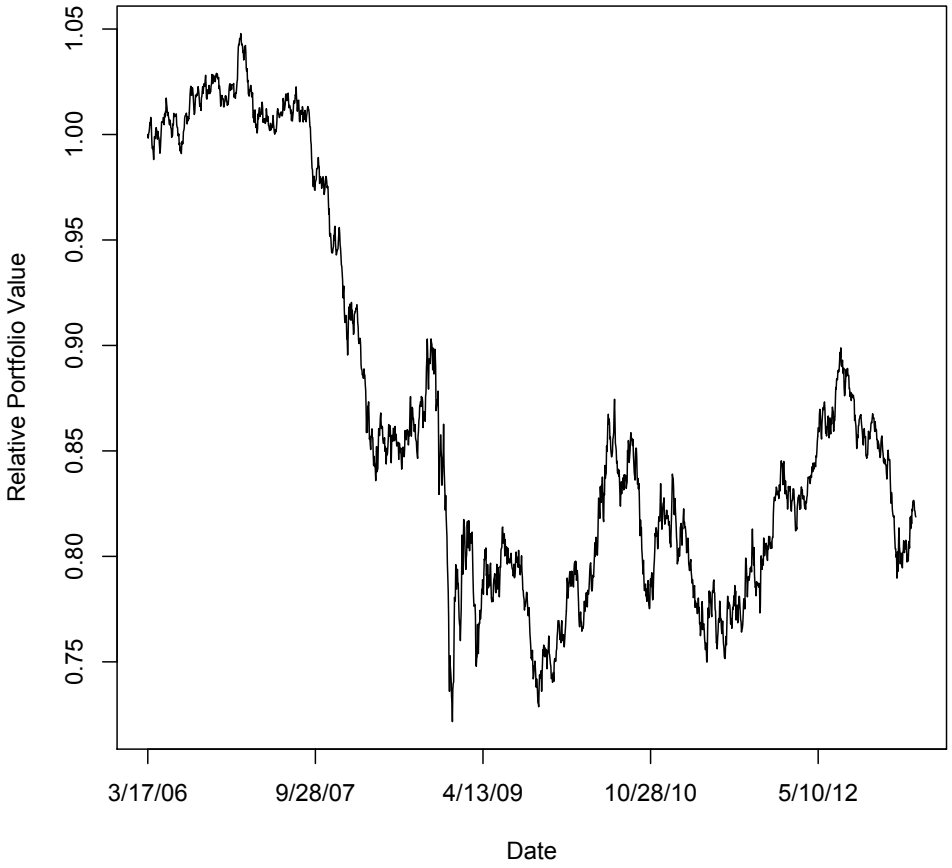

Figure 1: Relative Foreign Exchange Portfolio Value

they are all relatively simple and extremely computationally efficient, but differ in their underlying assumptions. The first two, Gaussian VaR and Nonparametric VaR, are commonly used in industry.

3.1 Gaussian Value at Risk

Gaussian VaR maintains the assumptions that returns for a given security or portfolio are stationary in their first two moments and normally distributed. If these assumptions are met, the sufficient statistics of the normal distribution can be estimated by historical returns. Let s_T be the sample standard deviation estimate using T past days of portfolio returns:

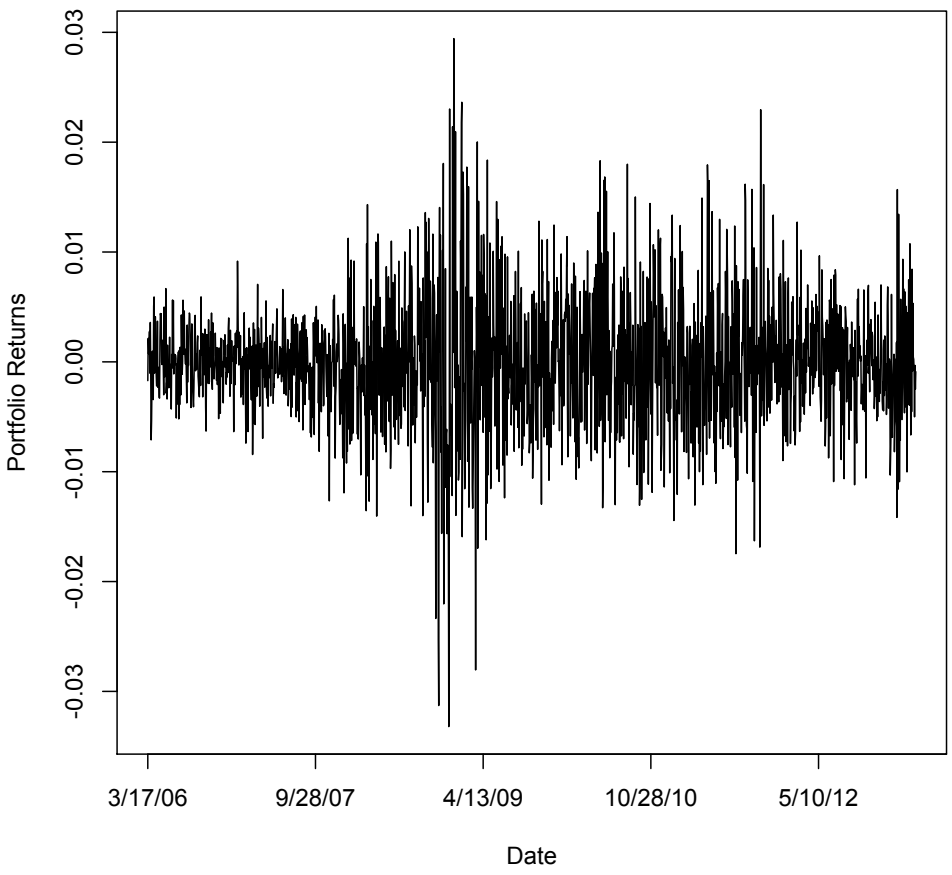

Figure 2: Daily Foreign Exchange Portfolio Returns

$$s_T(t) = \sqrt{\frac{1}{T-1}\sum_{k=t-1}^{t-T}(r_p - \bar{r_p})^2}. \tag{7}$$

Then, Gaussian VaR, denoted VaR^G, is the α quantile of the normal distribution with zero mean and standard deviation of s_T, which is calculated as

$$\text{VaR}^G(t) = \Phi^{-1}(\alpha|0,1) \cdot s_T(t), \tag{8}$$

where Φ^{-1} is the inverse normal cumulative distribution function. For the commonly calculated 95% VaR, equation (8) is approximately

$$\text{VaR}^G(t) \approx 1.64 \cdot s_T(t). \tag{9}$$

The size of the historical window used, T, is the only parameter one needs to specify (other than α) to compute Gaussian VaR. In this study, I use a historical window of 100 days. This length is long enough to obtain a precise estimate of realized volatility, while also allowing for temporal shifts in volatility that undoubtedly occur (see Figure 2 again). However, I do not claim that using 100 days is optimal and I comment on usage of different historical windows in the results section.

3.2 Nonparametric Value at Risk

Nonparametric VaR is popular at large financial institutions because it requires no distributional assumption and minimal computation. Let $\mathbf{R}$ be a vector of the past T days of portfolio returns, $\{r_p(t-1), r_p(t-2), \ldots, r_p(t-T)\}$, in ascending order (i.e. from most negative to most positive). Then, Nonparametric VaR, denoted VaR^{NP} is the α quantile of $\mathbf{R}$, or

$$\text{VaR}^{NP}(t) = \hat{F}^{-1}(\alpha), \tag{10}$$

where $\hat{F}^{-1}$ is the inverse empirical distribution function of $\mathbf{R}$. In the case of 95% VaR, and a historical window of 100 days, the Nonparametric VaR is simply the 5th worst return of the last 100 days.

3.3 Implied Volatility Value at Risk

The Implied Volatility VaR that I propose uses the at-the-money option implied volatilities found from traded (over-the-counter, in this case) option prices on the GBPUSD, EURUSD, and EURGBP exchange rates to estimate the future volatility and correlation of the portfolio's two assets (Walter and Lopez 2000). Firstly, the equation for the portfolio returns in equation (6) can be written more concisely as

$$r_p(t) = 2w(t) \cdot r_{GBP}(t) - 2(1 - w(t)) \cdot r_{EUR}(t), \tag{11}$$

where

$$w(t) = \frac{\text{GBPUSD(t-1)}}{\text{GBPUSD(1)}} \Big/ \left(\frac{\text{GBPUSD(t-1)}}{\text{GBPUSD(1)}} + \frac{\text{EURUSD(t-1)}}{\text{EURUSD(1)}} \right). \tag{12}$$

The variance of portfolio returns at time t, $\text{Var}(r_p(t))$, is

$$\begin{aligned}\text{Var}(r_p(t)) &= 4w(t)^2 \cdot \text{Var}(r_{GBP}(t)) + 4(1-w(t))^2 \cdot \text{Var}(r_{EUR}(t)) \\ &\quad -8w(t)(1-w(t)) \cdot \text{Cov}(r_{GBP}(t), r_{EUR}(t)). \end{aligned} \tag{13}$$

Since, the the weight $w(t)$ is actually known at time $t-1$, it is a constant at time t, making it a coefficient of the variance terms in equation (13). Denote the option implied volatilities at time t for the GBPUSD, EURUSD, and EURGBP exchange rates by $\sigma_{EUR}(t)$, $\sigma_{GBP}(t)$, and $\sigma_{E/G}(t)$, respectively. We can use squares of $\sigma_{EUR}(t)$ and $\sigma_{GBP}(t)$ for the GBPUSD and EURUSD variance terms in equation (13):

$$\begin{aligned}\text{Var}(r_p(t)) &= 4w(t)^2 \cdot \sigma_{GBP}(t)^2 + 4(1-w(t))^2 \cdot \sigma_{EUR}(t)^2 \\ &\quad -8w(t)(1-w(t)) \cdot \text{Cov}(r_{GBP}(t), r_{EUR}(t)). \end{aligned} \tag{14}$$

The last term in equation (14) can also be specified in terms of option-implied volatilities by using the option-implied covariance between GBPUSD and EURUSD. It is easy to show that the return on EURGBP, denoted $r_{E/G}$, is

$$r_{E/G}(t) = r_{EUR}(t) - r_{GBP}(t). \tag{15}$$

It follows that the variance of $r_{E/G}$ is

$$\text{Var}(r_{E/G}(t)) = \text{Var}(r_{EUR}(t)) + \text{Var}(r_{GBP}(t)) - 2\text{Cov}(r_{EUR}(t), r_{GBP}(t)), \tag{16}$$

which can be rearranged to give

$$\text{Cov}(r_{EUR}, r_{GBP}(t)) = 0.5\left[\text{Var}(r_{EUR}(t)) + \text{Var}(r_{GBP}(t)) - \text{Var}(r_{E/G}(t)\right]. \tag{17}$$

Using the appropriate option-implied volatilities, the option-implied covariance is

$$\text{Cov}(r_{EUR}, r_{GBP})(t) = 0.5(\sigma_{EUR}(t)^2 + \sigma_{GBP}(t)^2 - \sigma_{E/G}(t)^2). \tag{18}$$

Finally, combining equation (14) and equation (18), we can write the option-implied standard deviation of portfolio returns, σ_p as

$$\begin{aligned}\sigma_p(t) &= \big(4w(t)^2 \cdot \sigma_{GBP}(t)^2 + 4(1-w(t))^2 \cdot \sigma_{EUR}(t)^2 \\ &\quad -8w(t)(1-w(t)) \cdot 0.5(\sigma_{EUR}(t)^2 + \sigma_{GBP}(t)^2 - \sigma_{E/G}(t)^2)\big)^{\frac{1}{2}} . \end{aligned} \tag{19}$$

The Implied Volatility VaR, VaR^{IV}, is calculated by again assuming normality of returns:

$$\text{VaR}^{IV}(t) = \Phi^{-1}(\alpha|0,1) \cdot \sigma_p(t) \tag{20}$$

4 Evaluating Value at Risk Calculations

When comparing different VaR methodologies, it is imperative to specify criteria with which to evaluate their effectiveness. The difference between α and the observed event frequency is the first measure of efficacy, but is often the only metric of VaR's effectiveness considered (Berry). Additional information is needed to determine whether VaR is successfully characterizing risk. I propose that the following four criteria are important tests for VaR calculations:

1. VaR event frequency
2. VaR event autocorrelation
3. VaR event sizes
4. VaR volatility

In the following four sections I detail how we can measure and evaluate these four criteria.

4.1 Value at Risk Event Frequency

When attempting to calculate 95% VaR, say, an event rate of 10% should be very troubling, as VaR is being breached twice as often as expected. The holder of the portfolio will experience significant losses much more often than calculated and this is the first indication that this VaR calculation is failing. Of course, as time progresses the observed event frequency will change and seldom be exactly 5%, but an event rate close to this is important as a first order measure of calibration.

To quantify the difference between the observed and expected event frequency, we can model the number of VaR events with in a given time interval as being Poisson distributed, since VaR

events should be rare and independent (DeGroot and Schervish 2011, p 287). Let E be the number of events within T time periods. If events are independent (I discuss this in the next section), then the number of VaR events in a given time period are distributed as

$$E \sim \text{Pois}(0.05) = \frac{0.05^E e^{-0.05}}{E!}, \tag{21}$$

for the case of 95% VaR. The observed event rate, defined $\omega = E/T$, has a similar distribution:

$$\omega \sim \frac{1}{T} \frac{0.05^{\omega/T} e^{-0.05}}{(\omega/T)!}. \tag{22}$$

This distribution gives the probabilities of observed event frequencies given a true frequency of 5%. Thus, for an observed event frequency, $\omega_0 > 5\%$, we can take the area in the tail distribution in equation (22) to the right of ω_0 as a measure of statistical distance from ω_0 to 5%. This is the definition of a p-value, which being aware of its criticisms, I propose to use simply as a metric with which to compare the differences between observed and expected event frequencies across VaR methodologies.

4.2 Value at Risk Event Autocorrelation

Although an event rate of exactly 5% is surely a good sign, it does not mean a particular VaR methodology is capturing portfolio risk. If a VaR event occurs at time $t-1$, a correctly specified 95% VaR will still only have a 5% chance of being exceeded again at time t. If the underlying risk to the portfolio has increased, so should VaR. For an institution that holds capital against potential losses, VaR events that are highly autocorrelated will quickly deplete capital reserves and increase bankruptcy risk. This point is very important because a cluster of VaR events for an important institution can be systemically destabilizing.

One can test for autocorrelation by comparing the density for the number of time periods between VaR events with an exponential distribution. The exponential density is appropriate because independent VaR events should be Poisson-distributed and it is easy to show that Poisson waiting times are exponentially distributed. Let $\mathbf{\Lambda} = \{\lambda_1, \ldots, \lambda_n\}$ be the waiting times between an observed $n+1$ events. If the events were independent, the λ_i are distributed

$$\lambda_i \sim \omega e^{-\omega \lambda_i}, \quad \text{for } i = 1 \ldots n. \tag{23}$$

The observed event frequency ω is used instead of 5% because the λ_i need not be distributed exponentially with rate parameter 5% to imply that VaR events are independent. It may be

the case that events are independent, but occur more frequently that 5%. Divergence from the distribution in equation (23) is most appropriate to test for autocorrelation.

To assess how much a particular waiting time distribution diverges from an exponential distribution, I take the probability integral transform, which transforms the waiting times to a uniform distribution if and only if they are exponentially distributed. Doing so allows for easier comparison across different VaR calculations because all of the values will be between zero and one. Letting $\tilde{\mathbf{\Lambda}} = \{\tilde{\lambda_1}, \ldots, \tilde{\lambda}\}$ be the probability integral transformed $\mathbf{\Lambda}$, the λ_i are found as

$$\tilde{\lambda_i} = 1 - e^{-\omega\lambda_i}, \quad \text{for } i = 1 \ldots n. \tag{24}$$

I choose to compare the distribution of the $\tilde{\lambda_i}$ with the uniform distribution using quantile-quantile plots (Q-Q plots). As the name suggests, in a Q-Q plot each data point is the pair of values associated with a given quantile from each distribution. Having used the probability integral transform for the λ_i keeps the few, large waiting times from visually dominating these Q-Q plots. If the distributions are the same, the points in a Q-Q plot fall closely around the diagonal $y = x$ line. If the 95% confidence bands surrounding the data points do not mostly contain the diagonal line, this is evidence that the $\tilde{\lambda_i}$ are not uniformly distributed. In such a case, VaR waiting times are not exponentially distributed, implying that VaR events are not independent.

4.3 Value at Risk Event Sizes

An additional gauge of VaR's usefulness is the size of VaR shortfalls relative to VaR. If a bank calculates VaR to be \$1 million and proceeds to observe losses of \$3 million, this VaR calculation is not a helpful risk measure. This metric may defined as

$$\xi = \{r_p : r_p < 0, |r_p| > |\text{VaR}|\} / \text{VaR}, \tag{25}$$

which is shortfall divided by VaR. In the given example this quantity is $\xi = 3$. Admittedly, VaR is a quantile and does not attempt to describe shape the tail of the returns distribution, but in practice, the size of VaR events relative to VaR cannot be ignored. Intuitively, an extremely heavy tailed shortfall distribution indicates that VaR is often being fooled, in some sense, by large market moves.

4.4 Value at Risk Volatility

Fourthly, as changes in VaR may be taken as a signal to increase capital (or decrease exposure, leverage ect.) the volatility of VaR, itself, must be considered. A highly erratic VaR calculation can increase the transaction costs and costs of capital to its user. The volatility for a given VaR computation is simply the standard deviation of its log relatives.

5 Results

In this section I compare the Gaussian VaR, Nonparametric VaR, and Implied Volatility VaR along the four aforementioned criteria for a weekly, monthly and three-monthly VaR time horizon. Figures 3, 4, and 5 show the time series of the three VaR calculations and corresponding portfolio returns. In Figures 6 and 7, I show monthly VaR calculations using a 50-day and 200-day historical window for the Gaussian and Nonparametric VaR calculations. When using fewer past data, Gaussian and Nonparametric VaR adjust quicker, but also "forget" quicker, which leads to higher VaR event rates and volatility.

Table 1 shows the event frequencies for the three VaR methodologies and time horizons. At every horizon, using Gaussian or Nonparametric VaR results in markedly higher event frequencies than Implied Volatility VaR. As the horizon lengths are increased, the event frequencies increase substantially and can be taken to imply that none of the three methods is useful for forecasting VaR over a three-month period.

Table 1: 95% VaR Event Frequencies

	Gaussian VaR		Nonparametric VaR		IV VaR	
Time Horizon	Event Freq.	p-value	Event Freq.	p-value	Event Freq.	p-value
Weekly	7.6%	1.30E-02	6.5%	8.04E-02	5.7%	2.25E-01
Monthly	8.4%	2.42E-03	9.5%	1.73E-04	6.8%	5.35E-02
Three-Monthly	12.5%	1.64E-08	12.3%	4.25E-08	10.1%	3.97E-05

Figures 8, 9, and 10 show the Q-Q plots used to detect autocorrelation of VaR events at the weekly, monthly, and three-monthly time horizons, respectively. At the weekly horizon, the 95% confidence bands for all three VaR calculations mostly contain the diagonal line, though the

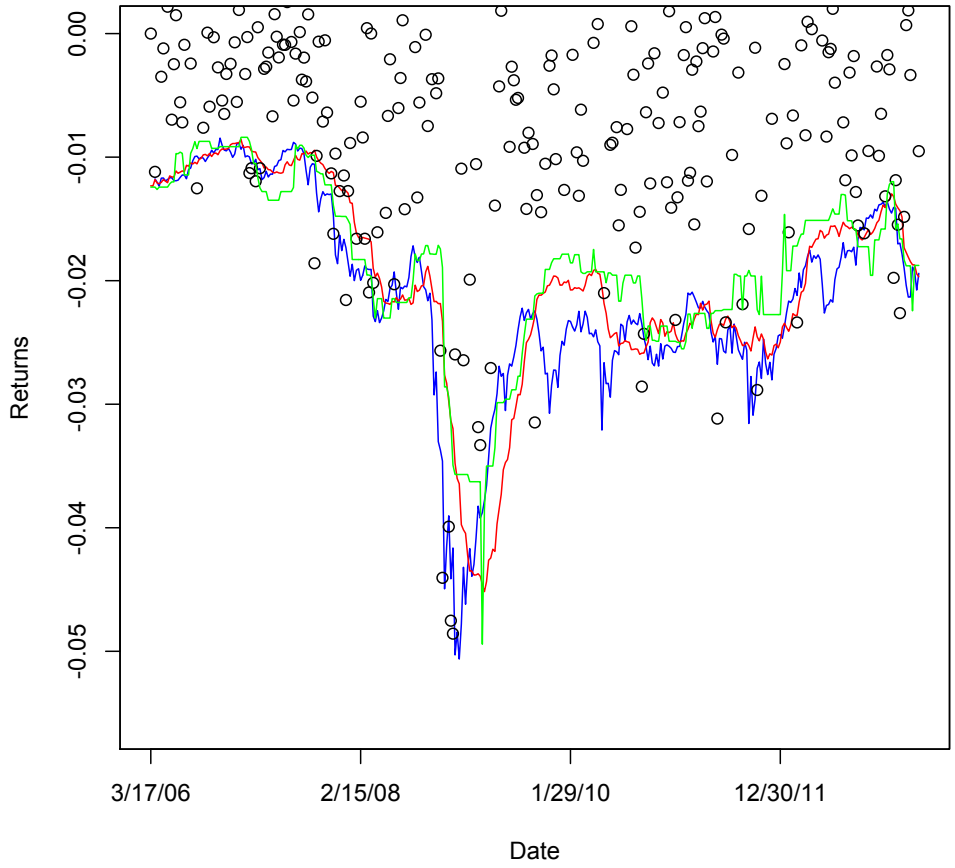

Figure 3: Value at Risk Time Series: Weekly Horizon

Gaussian VaR is the red line; Nonparametric VaR is the green line;

Implied Volatility VaR is the blue line; portfolio returns are the black dots.

Implied Volatility VaR points follow the diagonal much more closely. One may reasonably infer that significant autocorrelation is not found at the weekly horizon. At the monthly horizon, the 95% confidence bands mostly contain the diagonal in the Implied Volatility VaR plot, but not at all in the other two plots. The evidence from the Q-Q plots at the three-monthly horizon is consistent with the high empirical VaR frequencies, also suggesting that none of the models is producing reliable VaR estimates at the three month level.

The histograms for the quantity ξ defined in equation (25) are included for all VaR types and

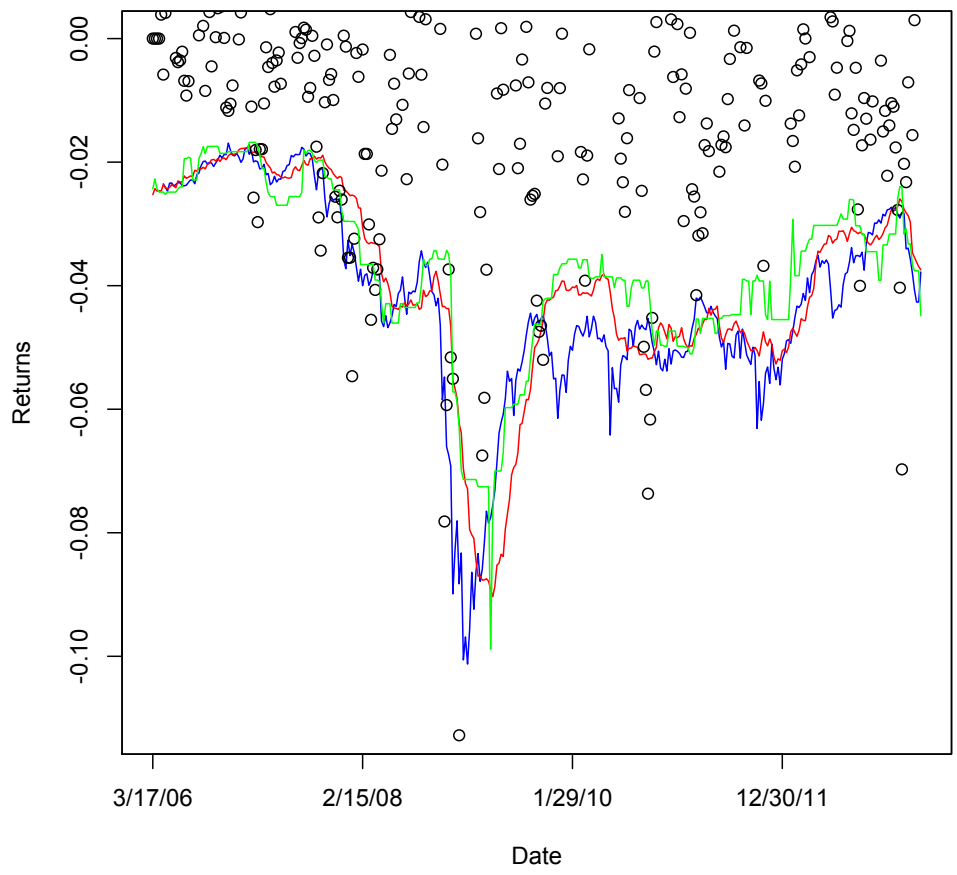

Figure 4: Value at Risk Time Series: Monthly Horizon
Gaussian VaR is the red line; Nonparametric VaR is the green line;
Implied Volatility VaR is the blue line; portfolio returns are the black dots.

time horizons in Figure 11, Figure 12, and Figure 13. A highly right-skewed distribution for ξ is undesirable because it indicates that when VaR events do happen, they are often much larger than VaR. Intuitively, a thinner-tailed distribution for ξ indicates that a particular VaR methodology is more aptly describing underlying risk. At all of the time horizons, using Gaussian or Nonparametric VaR yields a heavier tailed distribution than the Implied Volatility VaR equivalent. Additionally, the maximum value for ξ under Implied Volatility VaR is smaller than the largest value for the competing VaR methods.

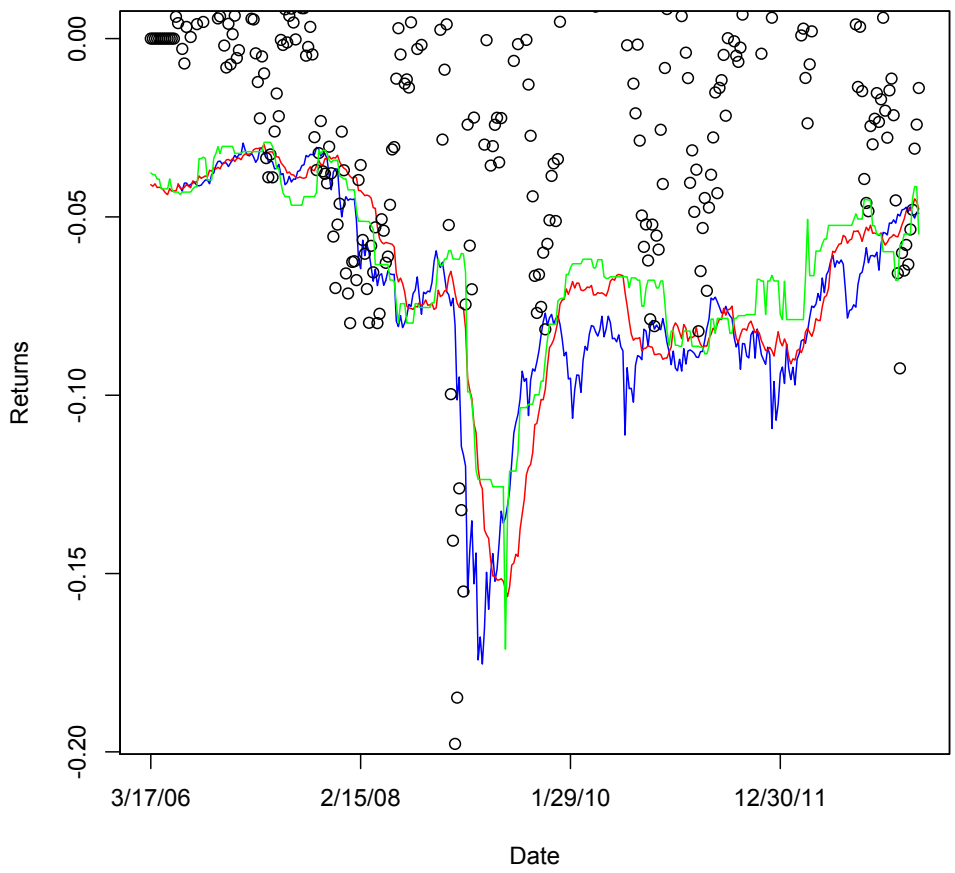

Figure 5: Value at Risk Time Series: Three-Monthly Horizon
Gaussian VaR is the red line; Nonparametric VaR is the green line;
Implied Volatility VaR is the blue line; portfolio returns are the black dots.

In Table 2, the daily volatilities for the VaR computations are given. The Gaussian VaR with volatility of around 3% is smaller than the rest because estimating volatility as the standard deviation of past returns is essentially an averaging process and should be expected to deliver smoother results. Although, all else equal, smaller VaR volatility is better, given the previous results, the lower volatility of Gaussian VaR illustrates its slow adjustment to changes in risk. Since Nonparametric VaR is often used in industry, that Implied Volatility VaR proved less volatile than Nonparametric VaR is encouraging for practical implementation.

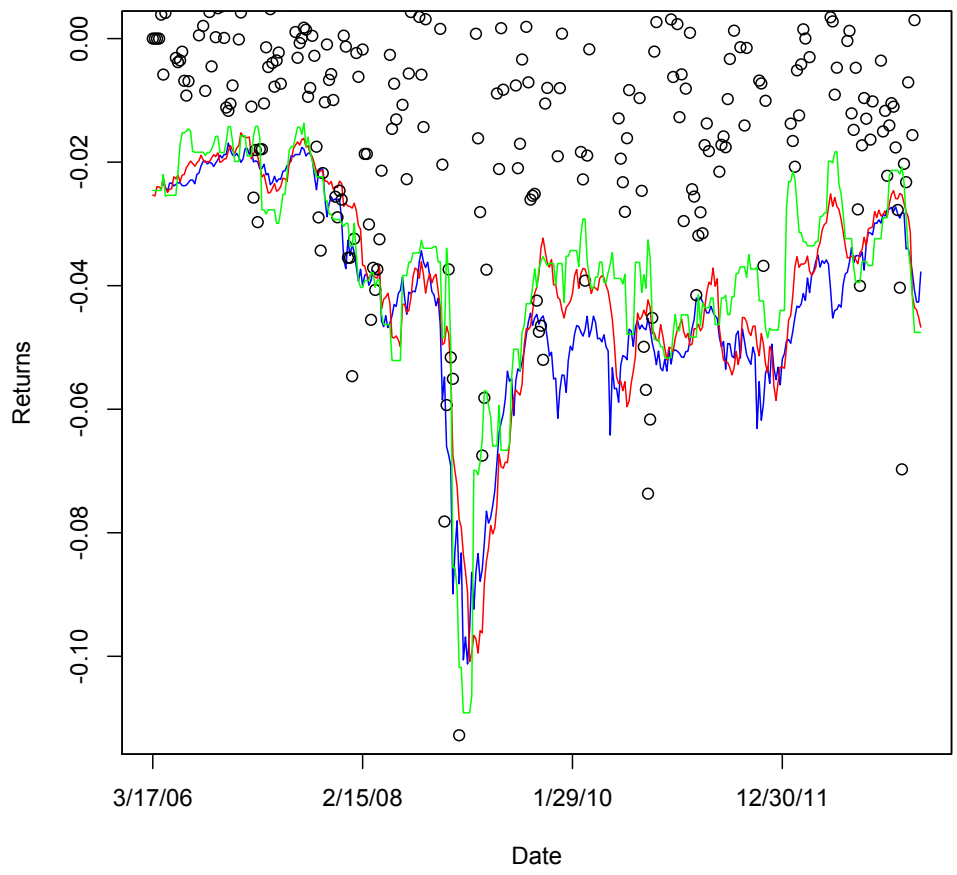

Figure 6: Value at Risk Time Series: Monthly-Monthly Horizon with 50-day Historical Window
Gaussian VaR is the red line; Nonparametric VaR is the green line;
Implied Volatility VaR is the blue line; portfolio returns are the black dots.

Table 2: 95% VaR Volatilities

Time Horizon	Gaussian VaR	Nonparametric VaR	IV VaR
Weekly	3.07%	6.44%	6.03%
Monthly	3.07%	6.39%	6.03%
Three-Monthly	2.97%	6.26%	5.96%

6 Conclusion

I have compared the effectiveness of three different VaR methodologies at three different time horizons for a foreign exchange portfolio that maintained a significant loss during the most recent

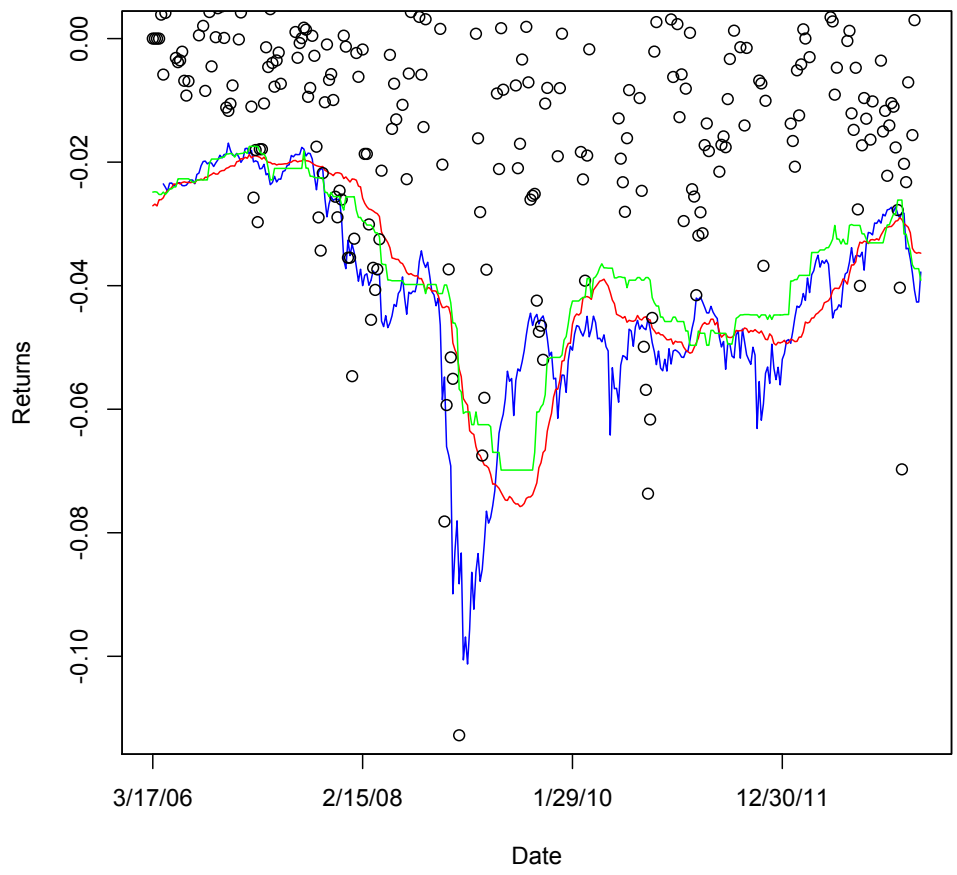

Figure 7: Value at Risk Time Series: Monthly-Monthly Horizon with 200-day Historical Window
Gaussian VaR is the red line; Nonparametric VaR is the green line;
Implied Volatility VaR is the blue line; portfolio returns are the black dots.

financial crisis. I chose event frequency, event autocorrelation, event size, and VaR volatility as important and relevant criteria to contrast the outcomes from implementing alternative VaR measures. I believe the results shown here to indicate that computing VaR using option-implied volatility is superior to the standard Gaussian or Nonparametric approaches. Using the market's forecast of volatility incorporates the beliefs and anticipations of informed market participants and, thus, adjusts at each time period to changes in risk. By stark contrast, the Gaussian and Nonparametric methodologies contain the assumption that the past market behavior is a good representation of

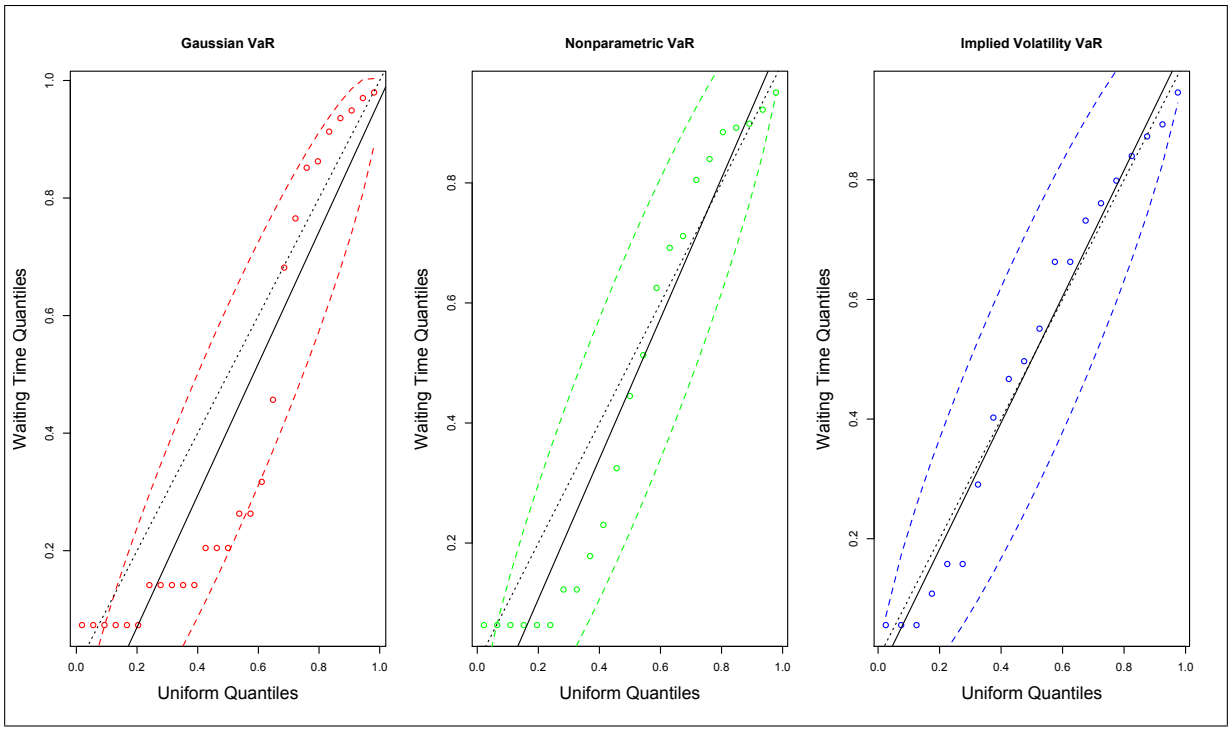

Figure 8: Waiting Times Quantile-Quantile Comparison with Uniform Distribution: Weekly Horizon

The dotted lines are y=x; the solid lines are regression lines; the dashed lines are 95% confidence bands

future activity. Given the dynamic and ever-evolving nature of markets, such a simple assumption is shown here to be a dangerous one. The results of this paper should be taken as evidence that using forward-looking sources of information may have considerable benefits to risk measurement and management.

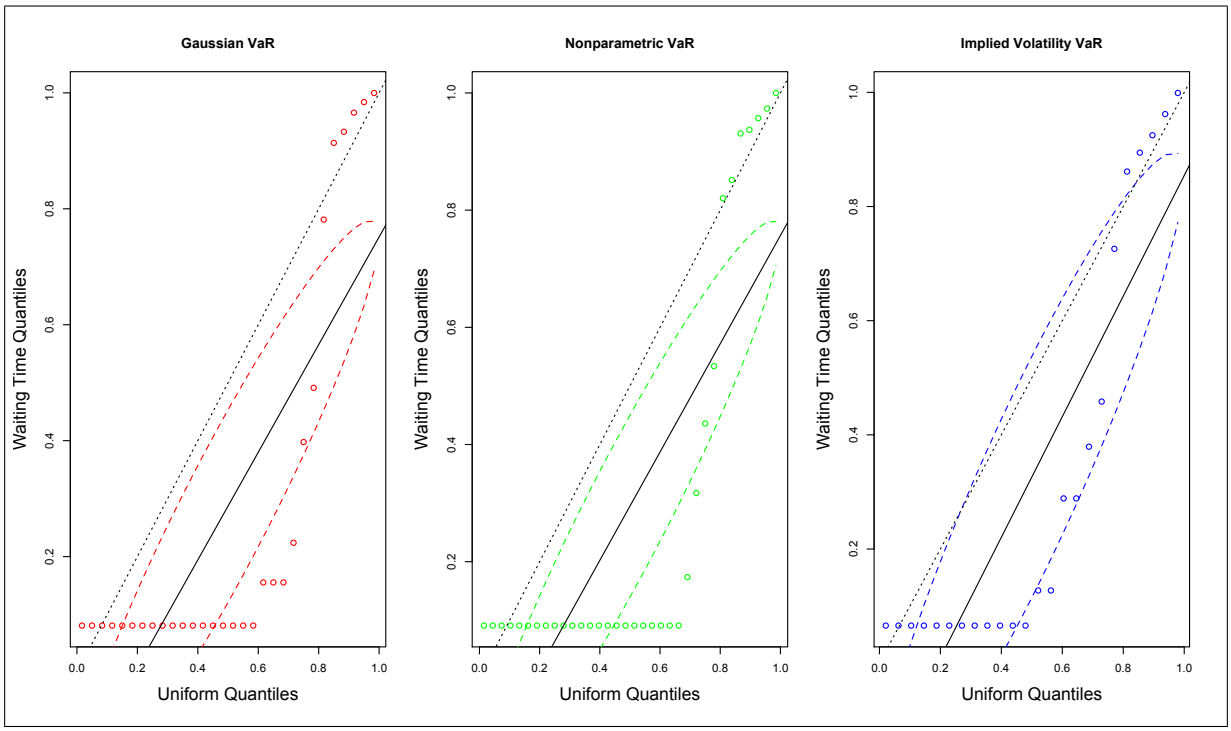

Figure 9: Waiting Times Quantile-Quantile Comparison with Uniform Distribution: Monthly Horizon

The dotted lines are y=x; the solid lines are regression lines; the dashed lines are 95% confidence bands

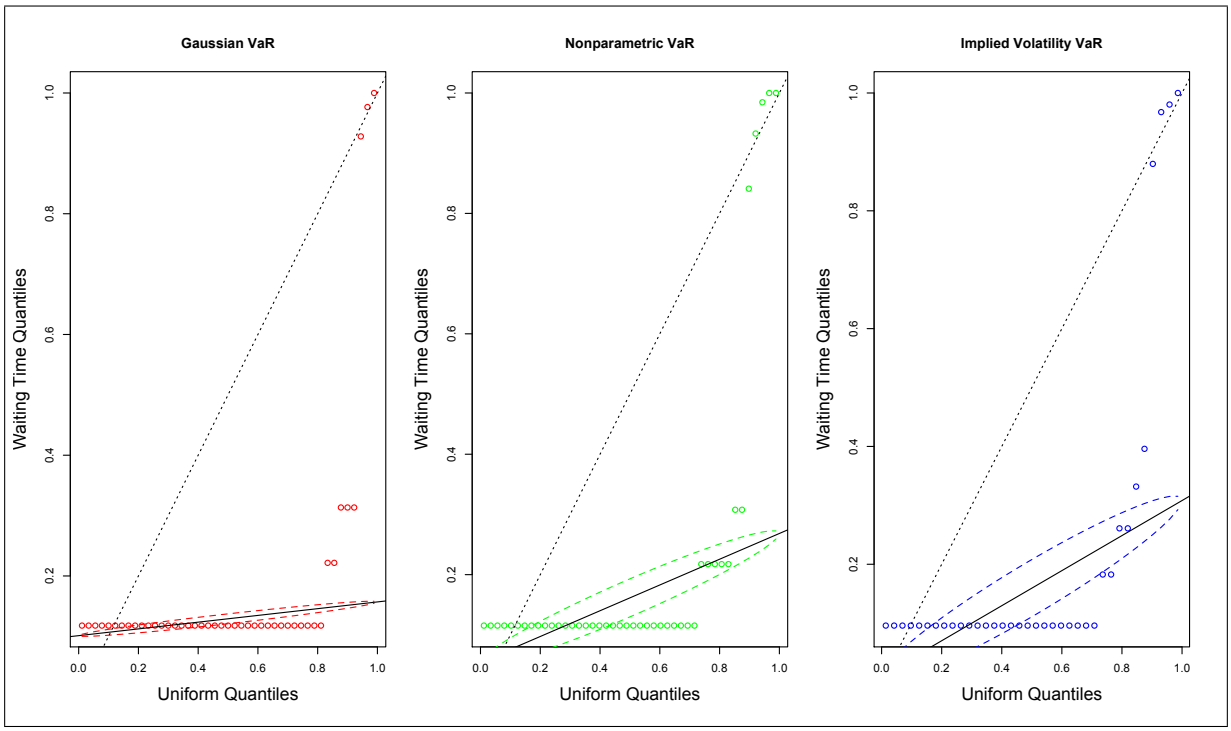

Figure 10: Waiting Times Quantile-Quantile Comparison with Uniform Distribution: Three-Monthly Horizon

The dotted lines are y=x; the solid lines are regression lines; the dashed lines are 95% confidence bands

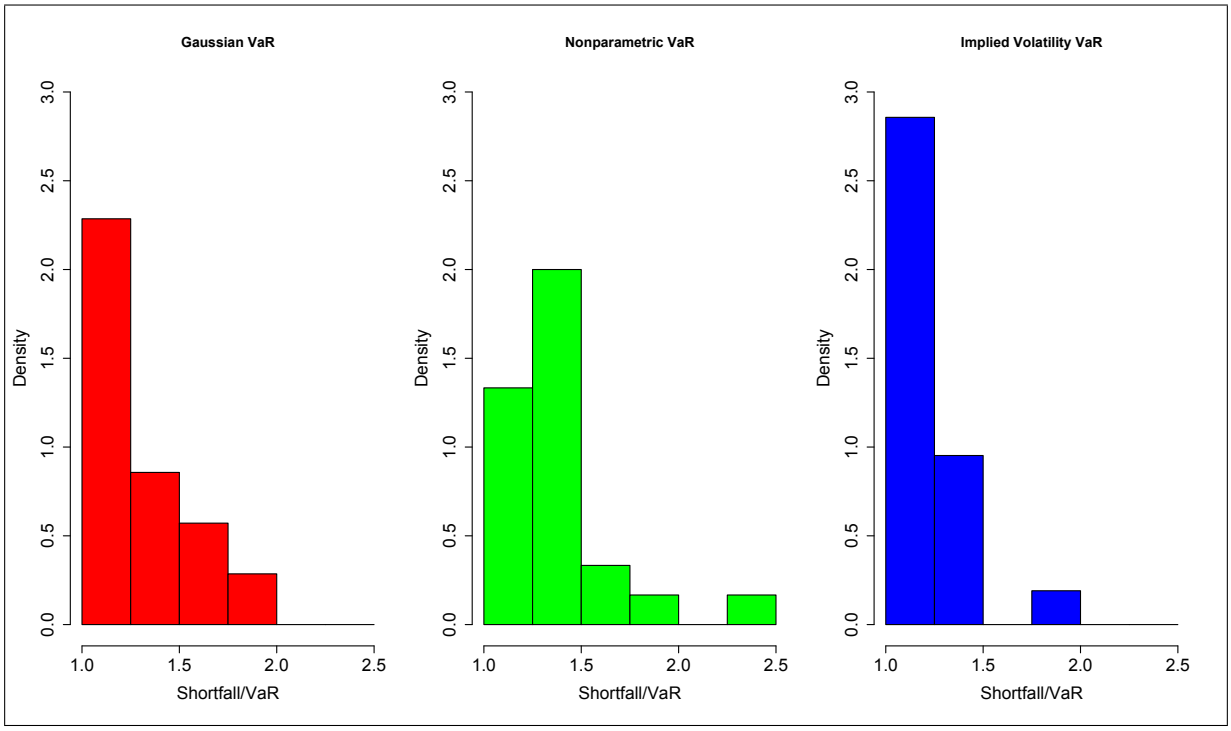

Figure 11: Shortfall/VaR Histogram: Weekly Horizon

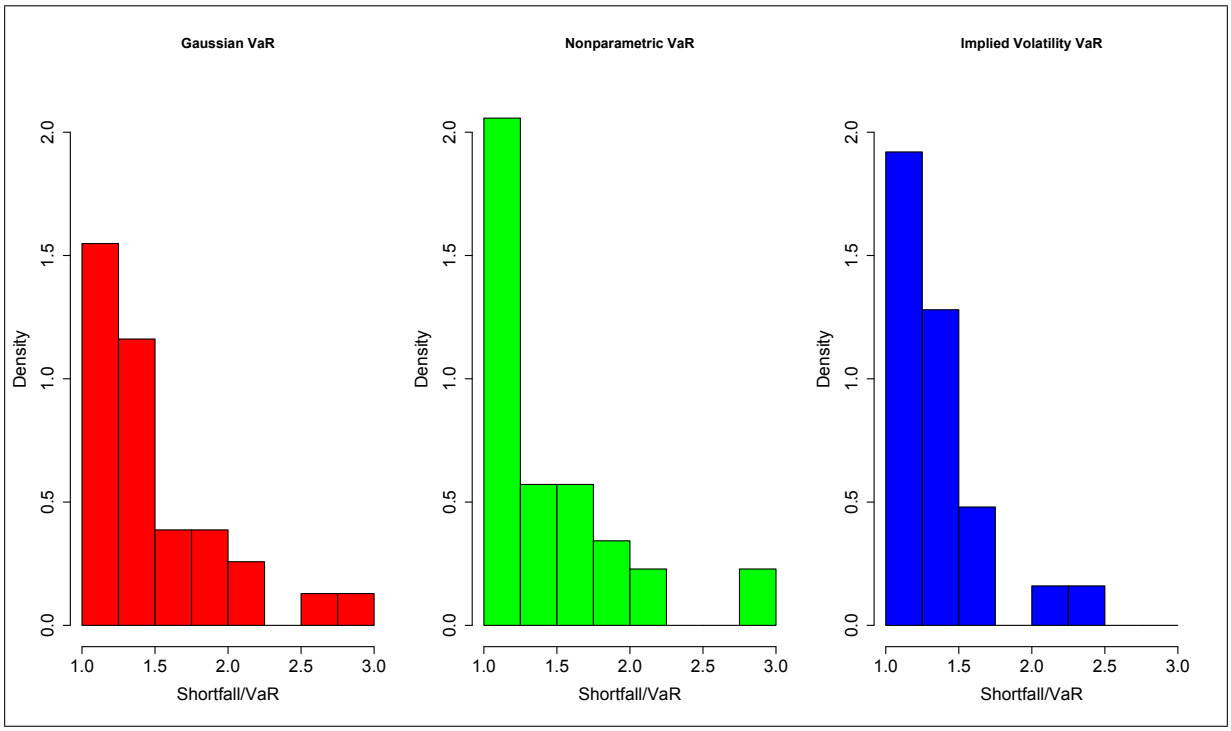

Figure 12: Shortfall/VaR Histogram: Monthly Horizon

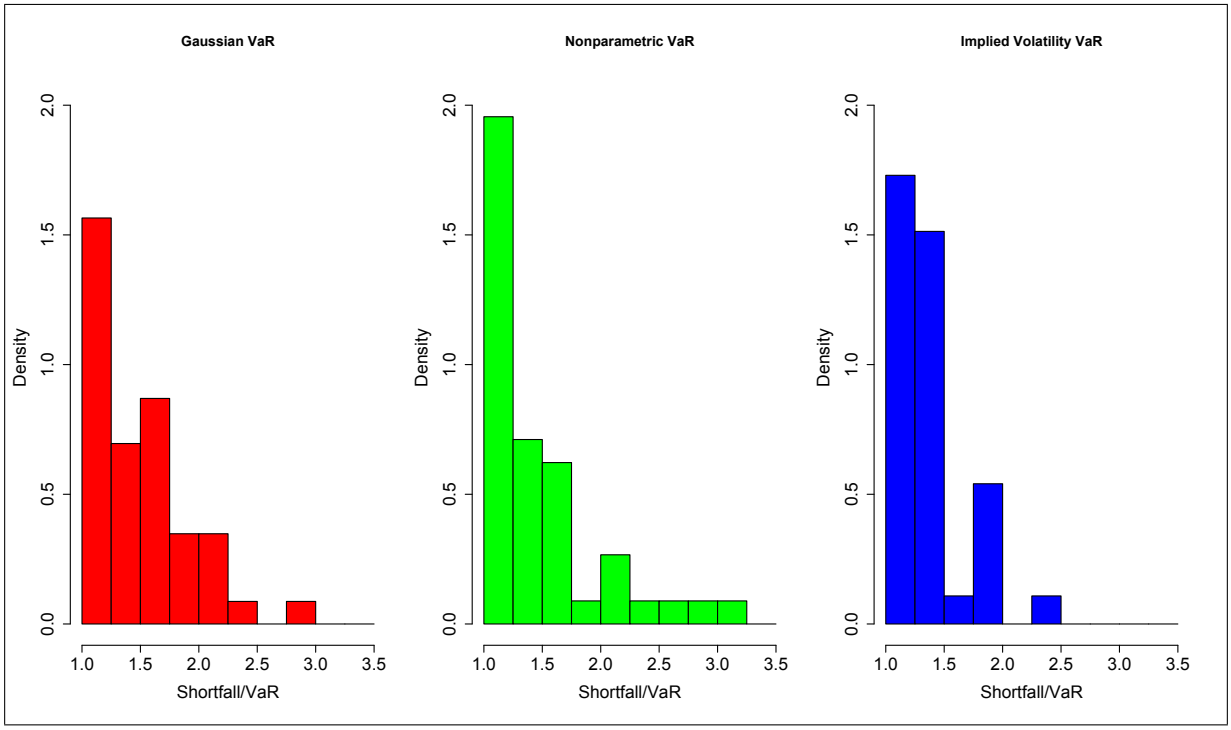

Figure 13: Shortfall/VaR Histogram: Three-Monthly Horizon

References

Romain Berry. Back testing value-at-risk. http://www.jpmorgan.com/tss/General/Back_Testing_Value-at-Risk/1159398587967.

Morris H. DeGroot and Mark J. Schervish. *Probability and Statistics (4th Edition)*. Pearson, 2011.

Paul Glasserman. *Monte Carlo Methods in Financial Engineering (Stochastic Modelling and Applied Probability) (v. 53)*. Springer, 2003.

Paul Wilmott. *Derivatives: the theory and practice of financial engineering*. Wiley Frontiers in Finance Series. J. Wiley, 1998.