

14 Introduction

15 This text provides additional methodological details related to synthetic snow wa-
 16 ter equivalent depth (SWE) observations (Section S1), synthetic hydropower production
 17 (Section S2), synthetic wholesale power prices (Section S3), the financial simulation model
 18 used to translate hydropower generation and power prices into net revenues (Section S4),
 19 and the method used to price index contracts (Section S5). Tables S1-S4 provide param-
 20 eter estimates for the models described in Sections S1-S3. Figures S1-S3 provide sup-
 21 port for Sections S1, S2, and S5 respectively, while Figures S4-S9 provide additional re-
 22 sults beyond the main text.

23 S1: Synthetic snow water equivalent depth (SWE)

24 Both February 1 and April 1 snow water equivalent depth (SWE) measurements
 25 were available for each year 1952-2016, except 1963, a total of 64 years. The historical
 26 SWE record is not found to exhibit any statistically significant trend in at an annual time
 27 step ($p = 0.45$ and 0.59 for February 1 and April 1, respectively), nor does it exhibit
 28 significant autocorrelation at an annual time step ($p > 0.05$ on Ljung-Box test for all
 29 lags up to 15 years for both February 1 and April 1 observations). Let \mathbf{s}_F and \mathbf{s}_A be the
 30 $n = 64$ years of February 1 and April 1 snow water equivalent depth (SWE) observa-
 31 tions described by the stationary random variables S_F and S_A . These two variables are
 32 inherently correlated through the seasonal snow accumulation process. It is important
 33 that any synthetic SWE observations capture the dependency structure of the original
 34 dataset.

35 Copulas provide a convenient way to capture this dependency. Let

$$36 \quad H(S_F, S_A) = C[F_F(S_F), F_A(S_A)] \quad (1)$$

37 where H is the joint cumulative distribution function (cdf) of S_F and S_A , F_F and F_A
 38 are their marginal cdf's, and C is the copula which dictates the dependency between the
 39 variables. Sklar showed that any joint distribution can be written in this form (Sklar,
 40 1973), which allows us to separate the treatment of the marginal distributions from their
 41 dependency structure.

The marginal distributions are found to be well described by gamma distributions, with probability density function (pdf)

$$f(S) = \frac{S^{k-1}e^{-x/\theta}}{\theta^k\Gamma(k)} \quad (2)$$

where $\Gamma()$ is the gamma function and k and θ are the shape and scale parameters, respectively. The parameter estimates \hat{k} and $\hat{\theta}$ are calculated independently for February and April using maximum likelihood estimation, and can be found in Table S1. We fail to reject the null hypothesis that the historical observations are drawn from the fitted gamma distribution using the Kolmogorov-Smirnov test of goodness of fit, for both February 1 ($p = 0.54$) and April 1 ($p = 0.98$) SWE.

Next, the copula describing the dependency structure between the two variables is estimated. Copula estimation consists of two steps: selection of a functional form and estimation of the functional parameters. A wide variety of functional forms have been studied in the literature (Frees & Valdez, 1998; Genest & Favre, 2007). We choose to use the Gaussian copula, part of the larger metaelliptical class of copulas, for a variety of reasons. Firstly, it is flexible and easily extensible to more than two variables, unlike other classes such as Archimedean copulas. Although the current work only uses two variables (S_F and S_A), we would like our workflow to be easily extensible to larger datasets in the future. Secondly, it is simple and allows for efficient generation of synthetic data. Lastly, it is shown to effectively capture the empirical dependency between February 1 and April 1 SWE, as will be shown shortly.

In Equation 1, H is a cumulative distribution function (cdf), taking on values from 0 to 1, as are the arguments to the copula ($F_F(S_F)$ and $F_A(S_A)$). Thus, the copula maps $(0, 1)^2 \rightarrow (0, 1)$. Rewriting the copula arguments as U_F and $U_A \in (0, 1)$, we can define the Gaussian copula

$$C(U_F, U_A) = \Phi_{2,\rho}(\Phi_1^{-1}(U_F), \Phi_1^{-1}(U_A)) \quad (3)$$

where Φ_1^{-1} is the inverse of the standard univariate normal cdf Φ_1 , and $\Phi_{2,\rho}$ is the bivariate joint cdf of standard normal variables with correlation ρ .

This shows that the copula only depends on the cumulative probabilities U , independently of the marginal distributions F . This allows for efficient sampling through the following algorithm (Wang, 1999): (1) Sample values (z_F, z_A) from the bivariate standard normal distribution with correlation ρ . For more than two variables, use the cor-

73 relation matrix Σ . (2) Calculate cumulative probabilities ($u_F = \Phi_1(z_F), u_A = \Phi_1(z_A)$).
 74 (3) Calculate the SWE values with equivalent cumulative probabilities under the gamma
 75 marginal distributions, ($s_F = F_F^{-1}(u_F), s_A = F_A^{-1}(u_A)$). This algorithm can be scaled
 76 up efficiently to generate many synthetic observations using modern statistical software.

77 The correlation is the only parameter that needs to be estimated for a Gaussian
 78 copula. Following Genest et al., (Genest, Favre, Béliveau, & Jacques, 2007), we estimate
 79 the correlation using Kendall's tau, a non-parametric analogue of correlation,

$$80 \quad \hat{\tau}_{FA} = \frac{C_{FA} - D_{FA}}{C_{FA} + D_{FA}} \quad (4)$$

81 where C_{FA} and D_{FA} are the number of concordant and discordant pairs, respectively,
 82 in the historical dataset. The pair (i, j) are defined as concordant if $(S_{F,i} - S_{F,j})(S_{A,i} -$
 83 $S_{A,j}) > 0$, and discordant otherwise. The estimate $\hat{\tau}$ is an asymptotically normal and
 84 unbiased estimator for the population value τ . Now the correlation estimate $\hat{\rho}$ can be
 85 calculated using the relation:

$$86 \quad \hat{\rho} = \sin\left(\frac{\pi \hat{\tau}}{2}\right) \quad (5)$$

87 The appropriateness of the Gaussian copula can be judged using a graphical ap-
 88 proach proposed by Genest & Favre (Genest & Favre, 2007). The approach is similar
 89 to the popular QQ-plot. Define the normalized order statistics,

$$90 \quad W_i = \frac{1}{n} \#\{j : S_{F,j} \leq S_{F,i}, S_{A,j} \leq S_{A,i}\} \quad (6)$$

91 In other words, W_i gives the fraction of observations with SWE values that are less than
 92 or equal to observation i in both February and April. W_i takes values on $(0, 1]$, and de-
 93 fines the empirical copula. The first step in the graphical approach is to calculate \mathbf{W}_{data} ,
 94 the $n = 64$ dimensional vector of normalized order statistics for the historical data. Then
 95 $M = 10,000$ synthetic samples of n years are generated from the fitted copula, and the
 96 order statistics ($\mathbf{W}_{fitted,1}, \dots, \mathbf{W}_{fitted,M}$) are calculated. Each \mathbf{W} vector is sorted from
 97 smallest to largest. The mean, 5th percentile, and 95th percentile are calculated for each
 98 of the $n = 64$ positions, yielding the vectors $\mathbf{W}_{fitted,mean}$, $\mathbf{W}_{fitted,p5}$, and $\mathbf{W}_{fitted,p95}$.
 99 Now the QQ-plot can be generated by scattering \mathbf{W}_{data} against $\mathbf{W}_{fitted,mean}$. If the ob-
 100 served data's dependency structure is well described by the fitted copula, then the plot-
 101 ted points should fall close to the one-to-one line. Sampling error bounds are estimated
 102 by $\mathbf{W}_{fitted,p5}$, and $\mathbf{W}_{fitted,p95}$. Figure S1 suggests that the fitted copula is a good fit for
 103 the historical SWE observations.

S2: Synthetic hydropower generation

In order to capture the relationship between snowpack and monthly hydropower generation, separate predictors are developed for each month of the water year, using the 29 water years available, 1988-2016.

For the first and last month of the water year (October and September), no statistically significant relationship (at a 10% significance level) is found between February 1 or April 1 SWE and hydropower generation. For these months expected generation is constant and the model can be written

$$G_{m,y}^{constant} = \beta_{0,m} + r_{m,y} \quad (7)$$

where $G_{m,y}^{constant}$ is the hydropower generation in month m of water year y , $\beta_{0,m}$ is the constant expected generation, and $r_{m,y}$ is the residual. The parameter estimate $\hat{\beta}_{0,m}$ is the sample mean.

The second class of model is a linear relationship between SWE and hydropower generation, written

$$G_{m,y}^{linear} = \beta_{0,m} + \beta_{1,m} S_{A/E,y} + r_{m,y} \quad (8)$$

where $S_{F/A,y}$ is either the February 1 or April 1 SWE value for water year y , and $\beta_{0,m}$ and $\beta_{1,m}$ are the intercept and slope parameters, respectively, to be estimated via linear regression. In cases where both February 1 and April 1 models are statistically significant (at a 10% significance level), the best model is selected by minimizing the Akaike information criterion (AIC). We find that February 1 SWE is the best linear predictor for November, December, and January, while April 1 SWE is used for February, July, and August.

For the peak snowmelt months of March through June, a clear upper threshold behavior is evident in the scatter plot of SWE vs hydropower generation. This reflects that fact that in the wettest years, some water may need to be spilled without generating hydropower. For this reason, generation is fit to April 1 SWE by minimizing the sum of squared residuals for the following piecewise linear model:

$$G_{m,y}^{piecewise} = \begin{cases} \beta_{0,m} + \beta_{1,m} S_{A,y} + r_{m,y} & , S_{A,y} < (\beta_{2,m} - \beta_{0,m})/\beta_{1,m} \\ \beta_{2,m} + r_{m,y} & , else \end{cases} \quad (9)$$

where $\beta_{0,m}$ and $\beta_{1,m}$ are the intercept and slope parameters of the increasing segment and $\beta_{2,m}$ is the expected generation in the constant segment above the SWE threshold.

134 The threshold is calculated based on these parameters as $(\beta_{2,m} - \beta_{0,m})/\beta_{1,m}$. Param-
 135 eter estimates for each monthly model described above can be found in Table S2, and
 136 models are visualized in Figure S2.

137 As described in the main text, the model residuals for the historical observations
 138 are deseasonalized based on month and SWE. The deseasonalized residuals are then fit
 139 to an autoregressive (AR) model in order to remove autocorrelation. Because all data
 140 points have been deseasonalized, the process has zero mean and thus no constant term.
 141 Only one and three month lags are found to be statistically significant (at the 5% level),
 142 resulting in the following linear regression:

$$143 \quad \tilde{r}_t = \varphi_1 \tilde{r}_{t-1} + \varphi_3 \tilde{r}_{t-3} + \varepsilon_t \quad (10)$$

144 where \tilde{r}_t is the deseasonalized residual for month t in the historical record, \tilde{r}_{t-1} and \tilde{r}_{t-3}
 145 are the deseasonalized residuals 1 and 3 months prior to observation t , and φ_1 , and φ_3
 146 are the regression parameters. The final model residuals ε_t are not found to exhibit sig-
 147 nificant autocorrelation (Ljung-Box test, $p > 0.18$ for all lags up to 36 months) and are
 148 not found to differ significantly from normality (Shapiro-Wilk test, $p = 0.79$). Param-
 149 eter estimates for the AR model can be found in Table S3.

150 **S3: Synthetic wholesale power prices**

151 The seven years of historical power prices are first log-transformed and deseason-
 152 alized, as described in the main text. Let x_t be the deseasonalized log-prices ($\tilde{p}_{m,y}$ in the
 153 main text) for the month t in the time series corresponding to (m, y) . Now the desea-
 154 sonalized log prices can be fit to a seasonal autoregressive moving average (SARMA) model.
 155 Let $\text{SARMA}(p, q)(P, Q)_{12}$ represent a SARMA model that combines an autoregressive
 156 model with lags of $(1, \dots, p)$ months and $12*(1, \dots, P)$ months, and a moving average er-
 157 ror model with lags of $(1, \dots, q)$ and $12*(1, \dots, Q)$ months. Because all data points have
 158 been deseasonalized, the process has zero mean and thus no constant term. This model
 159 can be written:

$$160 \quad x_t = \sum_{i=1}^p \varphi_i x_{t-i} + \sum_{j=1}^P \varphi_{12j} x_{t-12j} + \sum_{k=1}^q \theta_k \varepsilon_{t-k} + \sum_{l=1}^Q \theta_{12l} \varepsilon_{t-12l} + \varepsilon_t \quad (11)$$

161 where φ_L is the autoregressive parameter for lag L , θ_L is the moving average error pa-
 162 rameter for lag L , and ε_t is the prediction error for month t . In order to choose the ap-
 163 propriate model order, regression parameters are estimated for all combinations of (p, q, P, Q) ,

164 where each parameter is either 0 or 1 (a total of 16 models). From among the four mod-
 165 els with similarly low Bayesian Information Criterion (BIC) values ($BIC_{(1,1,0,1)} = 123.35$,
 166 $BIC_{(1,1,1,0)} = 123.90$, $BIC_{(1,0,0,1)} = 125.15$, $BIC_{(1,0,1,0)} = 125.71$), the SARMA(1, 0), (0, 1)₁₂
 167 model is selected as the model which best matches the monthly patterns of the histor-
 168 ical data. This model consists of a single lag of one month for the autoregressive model,
 169 plus a moving average error model with a single lag of twelve months, and can be writ-
 170 ten:

$$171 \quad x_t = \varphi_1 x_{t-1} + \theta_{12} \varepsilon_{t-12} + \varepsilon_t \quad (12)$$

172 Parameter estimates for the SARMA model can be found in Table S4. The resid-
 173 uals from the SARMA model are not found to exhibit significant autocorrelation ($p >$
 174 0.1 on Ljung-Box test for all lags up to 36 months) and are not found to deviate signif-
 175 icantly from a normal distribution (Shapiro-Wilk test, $p = 0.51$).

176 **S4: Revenue model**

177 Given synthetic hydropower generation G_t and wholesale power prices P_t , the sim-
 178 ulated revenues R_t can be written:

$$179 \quad R_t = r^M D^M + r^I D_t^I + P_t (G_t - D^M - D_t^I) \quad (13)$$

180 where r^M and r^I are the fixed volumetric rates for municipal customers and irrigation
 181 districts, respectively, and D^M and D_t^I are the fixed municipal demand and variable ir-
 182 rigation district demand.

183 The municipal rate ($r^M = \$0.1049/\text{kWh}$) and demand ($D^M = 998.405 \text{ GWh/year}$)
 184 and are taken as their 2016 values from public financial statements (San Francisco Pub-
 185 lic Utilities Commission, 2016). The demand is calculated as the sum of four classes: “Gen-
 186 eral Fund Rate Subsidized,” “Enterprise Rate,” “Non-city Agencies,” and “Moccasin/City
 187 of Riverbank,” and the fixed rate is calculated as a sales-weighted average of the four
 188 class rates. Given that heating and cooling tend to drive variability in electricity demand,
 189 retail demand is fairly constant throughout the year due to the mild Bay Area climate
 190 in which temperatures do not deviate widely on a seasonal basis (San Francisco Public
 191 Utilities Commission, 2015). In light of this, it is assumed for this work that the demand
 192 is spread evenly across the year.

193 The fixed rate for Modesto and Turlock Irrigation Districts (MTID) is taken as the
 194 sales-weighted average value of the 2016 rates for the two districts in the financial state-
 195 ment (San Francisco Public Utilities Commission, 2016). In this work, we assume that
 196 MTID purchases a fixed fraction of excess hydropower generation above municipal de-
 197 mand, so that demand is fit to the model:

$$198 \quad D_t^I = c \max(G_t - D^M, 0) \delta_{P_t > r_I} \quad (14)$$

199 where c , the fraction of excess hydropower generation that MTID will purchase, is a pa-
 200 rameter to be fit using linear regression. $\delta_{P_t > r_I}$ is an indicator function ensuring that
 201 MTID only purchases power from the utility when the fixed rate is lower than the whole-
 202 sale power rate. Using the historical annual sales volumes from 2010-2016 (San Francisco
 203 Public Utilities Commission, 2016), we find $c = 0.483$ ($p < 0.01$). This means that
 204 48.3% of the excess power in any given month is sold to MTID. The other 51.7% of ex-
 205 cess power is sold into the wholesale market.

206 **S5: Index contract pricing**

207 In order to find the contract loading and the corresponding reference value for the
 208 SWE index, which separates positive payouts at low SWE values and negative payouts
 209 at high SWE values, the capped contract for differences (CFD) structure is decomposed
 210 into two parts. The CFD can be seen as the sum of a “long put” position and a “short
 211 capped call” position, as shown in Supporting Information Figure S3. The long put po-
 212 sition provides linearly increasing payouts when the SWE index falls below a threshold
 213 called the “strike”, in exchange for an annual premium, while the short call position re-
 214 quires making payments when the SWE index falls below the strike, in return for receiv-
 215 ing an annual premium. A cap on call payments is applied at some higher threshold. The
 216 payouts shown in Figure S3 have been netted of premiums. The strike for both contracts
 217 is set at the 50th percentile of the SWE distribution (24.53 inches), while the cap is set
 218 at the 95th percentile (48.44 inches). The swap contract is created by summing these
 219 two positions.

220 In order to price the premiums, the Wang transform, as described in Section 2.6.3
 221 of the main text, can be performed numerically for any payout cumulative distribution
 222 function (cdf) $F(x)$. In this case, $F(x)$ is the empirical cdf built from the synthetic con-
 223 tract payout time series, which is itself a function of the synthetic SWE index time se-

224 ries. For example, for the put contract,

$$225 \quad x = V * \max(k - s, 0) \quad (15)$$

226 where V is the contract slope in dollars per inch, k is the strike value set as the median
 227 of the SWE distribution, and s is the SWE value observed. The contract premium is then
 228 calculated as the expected value of x under the risk-adjusted probability density func-
 229 tion (pdf) $f^*(x)$. Thus, the expected net payout, x minus the premium, is equal to zero
 230 under the risk-adjusted pdf. After calculating the premiums for the long put and short
 231 capped call, the CFD structure is equivalent to the sum of these two positions. The “load-
 232 ing” in Section 2.6.3 of the main text is equal to the premium for the long put position
 233 minus the premium for the short capped call position. The reference value of SWE for
 234 the CFD is the intersection of net payout function with the x-axis.

235 Following other work on weather derivative contracts (Baum, Characklis, & Serre,
 236 2018; Foster, Kern, & Characklis, 2015; Wang, 2002), the market price of risk param-
 237 eter is set as $\lambda = 0.25$ for the baseline value for the put contract in this study. How-
 238 ever, values between $\lambda = 0$ (no loading) and 0.5 (high loading) are included in the sen-
 239 sitivity analysis, as described in Section 2.9 of the main text, and the effect of this pa-
 240 rameter on the contract structure can be seen in Figure 5 of the main text. Because the
 241 short capped call position is used in payment for the long put position, rather than a sep-
 242 arate hedging contract, it is priced in the actuarially fair manner (i.e., $\lambda = 0$). If a risk
 243 loading were applied to this position, this would imply that the contract seller is will-
 244 ing to accept a lower (expected value) premium if it comes in variable payments rather
 245 than fixed payments, which is unlikely in practice.

246 **References**

- 247 Baum, R., Characklis, G. W., & Serre, M. L. (2018). Effects of Geographic Diversifi-
 248 cation on Risk Pooling to Mitigate Drought-Related Financial Losses for Water
 249 Utilities. *Water Resources Research*, *54*, 1–19. doi: 10.1002/2017WR021468
- 250 Foster, B. T., Kern, J. D., & Characklis, G. W. (2015). Mitigating hydrologic fi-
 251 nancial risk in hydropower generation using index-based financial instruments.
 252 *Water Resources and Economics*, *10*, 45–67. doi: 10.1016/j.wre.2015.04.001
- 253 Frees, E. W., & Valdez, E. A. (1998). Understanding relationships using copu-
 254 las. *North American Actuarial Journal*, *2*(1), 1–25. doi: 10.1080/10920277

- 255 .1998.10595667
- 256 Genest, C., & Favre, A.-C. (2007). Everything You Always Wanted to Know about
 257 Copula Modeling but Were Afraid to Ask. *Journal of Hydrologic Engineering*,
 258 12(4), 347–368. doi: 10.1061/(asce)1084-0699(2007)12:4(347)
- 259 Genest, C., Favre, A. C., Béliveau, J., & Jacques, C. (2007). Metaelliptical copu-
 260 las and their use in frequency analysis of multivariate hydrological data. *Water*
 261 *Resources Research*, 43(9), 1–12. doi: 10.1029/2006WR005275
- 262 San Francisco Public Utilities Commission. (2015). *Hetch Hetchy Water & Power,*
 263 *Financial Statements, June 30, 2015 and 2014* (Tech. Rep.).
- 264 San Francisco Public Utilities Commission. (2016). *Comprehensive Annual Financial*
 265 *Report, Fiscal Years Ended June 30, 2016 and 2015* (Tech. Rep.).
- 266 Sklar, A. (1973). Random Variables, Joint Distribution Functions, and Copulas. *Ky-*
 267 *bernetika*, 9(6), 449–460.
- 268 Wang, S. S. (1999). “Understanding Relationships Using Copulas,” Edward Frees
 269 and Emiliano Valdez, January 1998. *North American Actuarial Journal*, 3(1),
 270 137–142. doi: 10.1080/10920277.1999.10595785
- 271 Wang, S. S. (2002). A Universal Framework for Pricing Financial and Insurance
 272 Risks. *ASTIN BULLETIN*, 32(2), 213–234. doi: 10.2143/AST.32.2.1027

Tables**Table S1.** Parameter estimates for synthetic snow water equivalent (SWE) model. All parameters unitless.

Parameter	Date	Symbol	Estimate
Gamma shape	February 1	\hat{k}_F	2.9167
	April 1	\hat{k}_A	5.1693
Gamma scale	February 1	$\hat{\theta}_F$	6.4355
	April 1	$\hat{\theta}_A$	5.7539
Copula correlation	NA	$\hat{\rho}_{FA}$	0.8017

Table S2. Parameter estimates for monthly models of synthetic hydropower generation based on SWE. Units for $\hat{\beta}_{0,m}$ and $\hat{\beta}_{2,m}$ are GWh per month, and units for $\hat{\beta}_{1,m}$ are GWh per month per inch of SWE.

Month	SWE predictor	Model	$\hat{\beta}_{0,m}$	$\hat{\beta}_{1,m}$	$\hat{\beta}_{2,m}$
October	NA	Constant	83.76	NA	NA
November	Feb. 1	Linear	59.29	0.9182	NA
December	Feb. 1	Linear	36.97	3.251	NA
January	Feb. 1	Linear	43.05	3.639	NA
February	Apr. 1	Linear	14.57	3.942	NA
March	Apr. 1	Piecewise	48.67	4.691	230.8
April	Apr. 1	Piecewise	91.66	3.774	226.4
May	Apr. 1	Piecewise	126.8	3.179	235.9
June	Apr. 1	Piecewise	33.29	6.045	232.8
July	Apr. 1	Linear	45.54	3.532	NA
August	Apr. 1	Linear	89.94	0.8156	NA
September	NA	Constant	93.68	NA	NA

Table S3. Parameter estimates for autoregressive model for synthetic hydropower generation residuals. All parameters unitless.

Parameter	Symbol	Estimate
Lag-1 weight	$\hat{\varphi}_1$	0.5405
Lag-3 weight	$\hat{\varphi}_3$	-0.1203

Table S4. Parameter estimates for seasonal autoregressive moving average (SARMA) model for deseasonalized log power prices. All parameters unitless.

Parameter	Symbol	Estimate
Lag-1 autoregression weight	$\hat{\varphi}_1$	0.8663
Lag-12 moving average error weight	$\hat{\theta}_{12}$	-0.3740

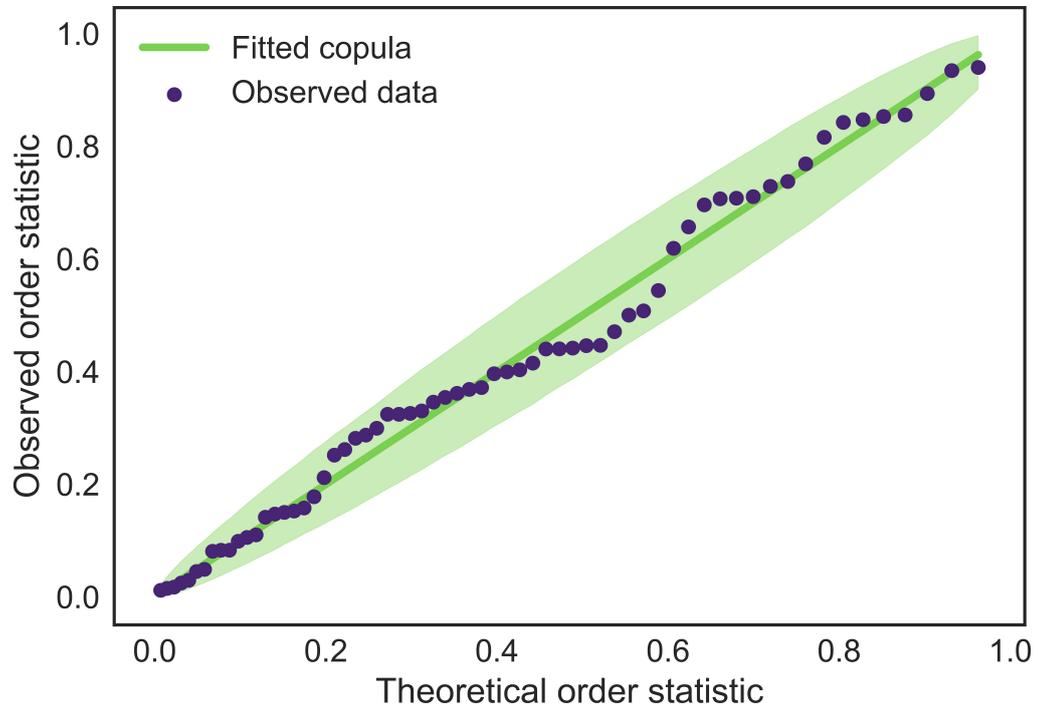
Figures

Figure S1. QQ-plot showing fit between normalized order statistics of the observed data and the fitted Gaussian copula. Shaded region shows the 5th and 95th percentile sampling error for the Gaussian copula.

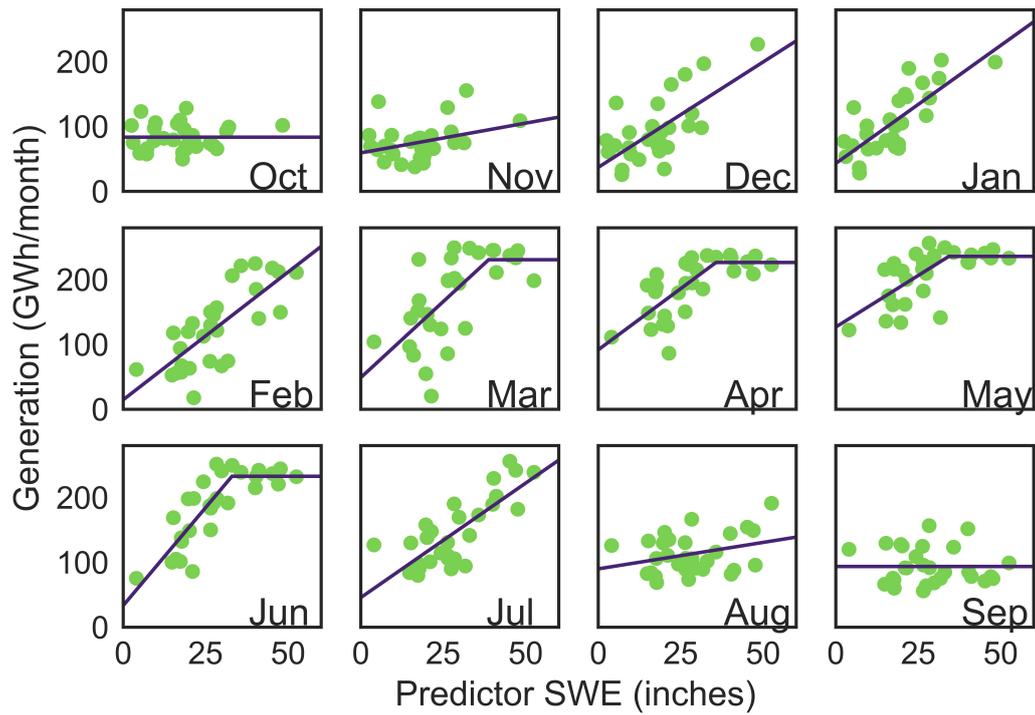


Figure S2. Hydropower generation as a function of snow water equivalent depth (SWE) for each month in the water year. Fitted models (purple lines) shown against historical data (green dots). The x-axis “predictor SWE” is February 1 SWE for October through January, and April 1 SWE for February through September.

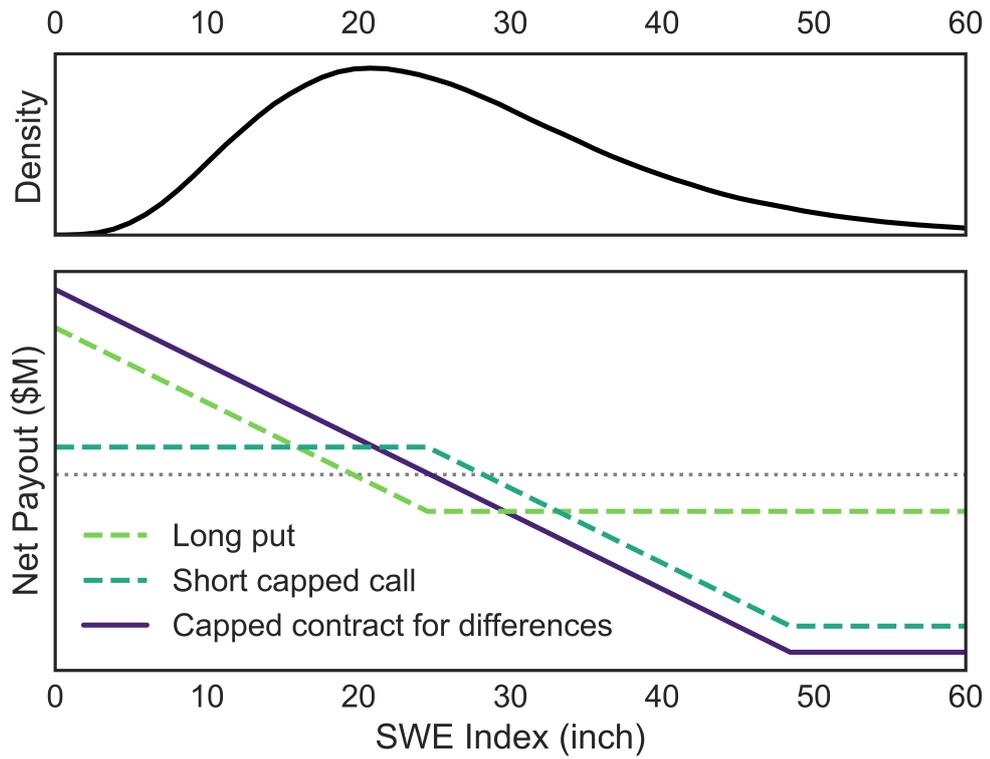


Figure S3. (top) Probability density for SWE index, a weighted average of February 1 and April 1 observations. (bottom) Net payouts for three possible contract structures: a long put position with a strike at the 50th percentile (24.53 inch), a short call position with a strike at the 50th percentile and a cap at the 95th percentile (48.44 inch), and a capped contract for differences (CFD) equal to the sum of the two previous positions.

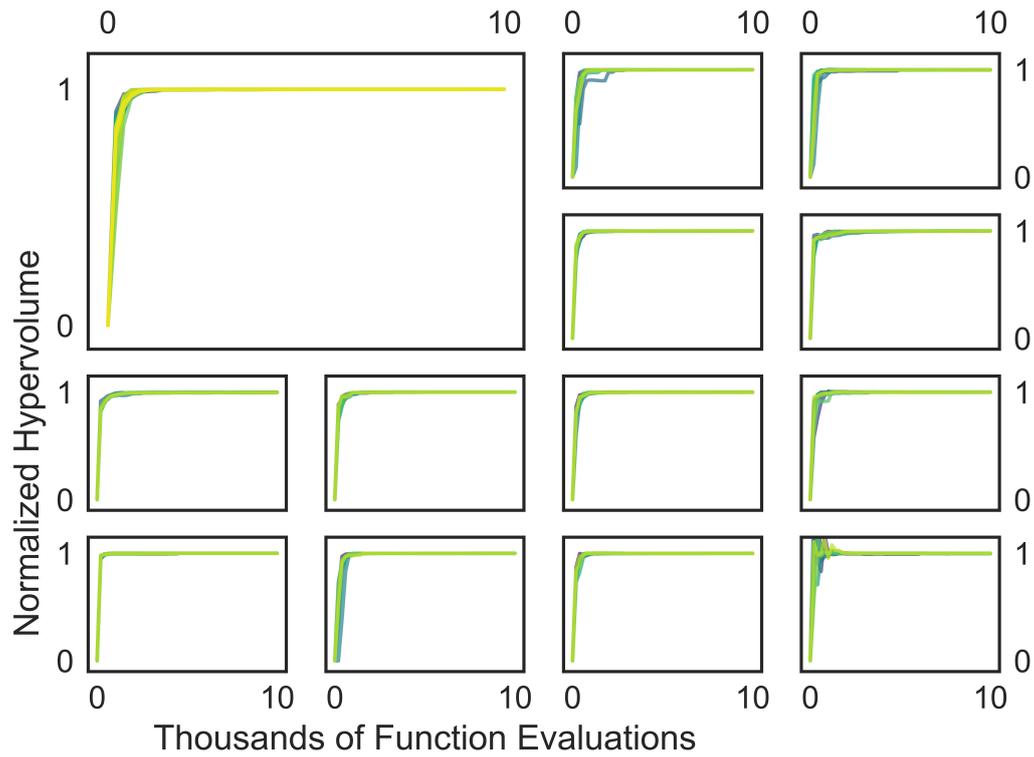


Figure S4. (top left) Normalized hypervolume of approximate Pareto sets from Borg MOEA with 50 random seeds for baseline state of world (SOW). (all others) Normalized hypervolume of approximate Pareto sets from Borg MOEA with 10 random seeds each for 12 randomly selected SOWs.

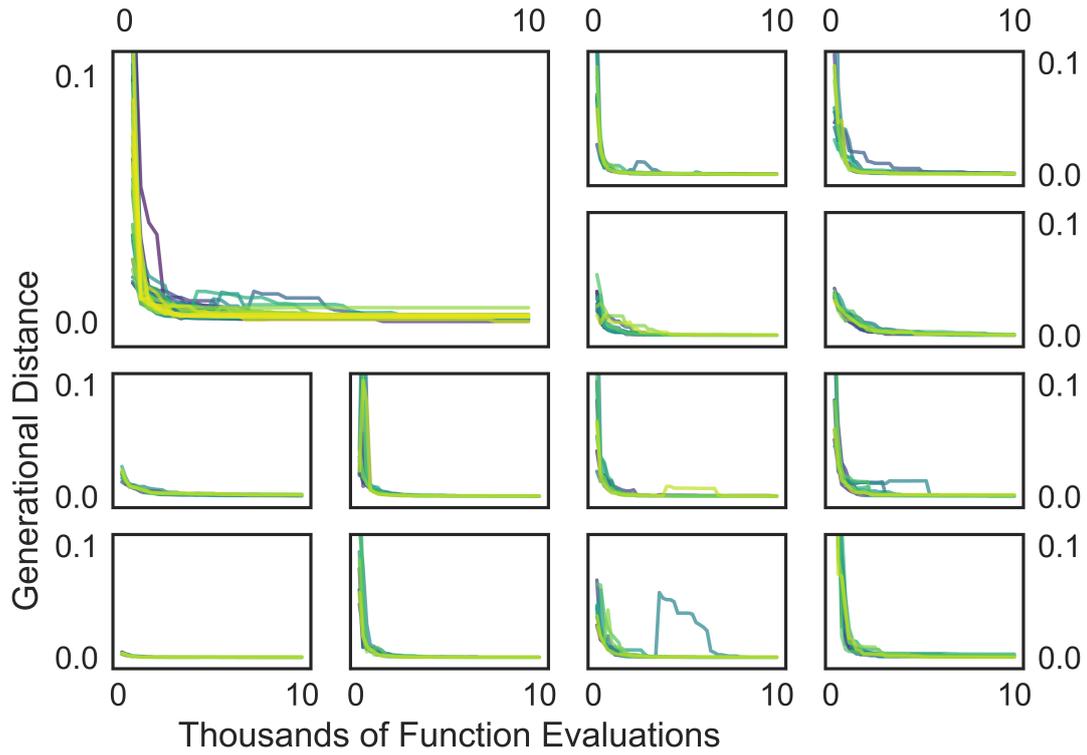


Figure S5. (top left) Generational distance metric for approximate Pareto sets from Borg MOEA with 50 random seeds for baseline state of the world (SOW). (all others) Generational distance metric for approximate Pareto sets from Borg MOEA with 10 random seeds each for 12 randomly selected SOWs.

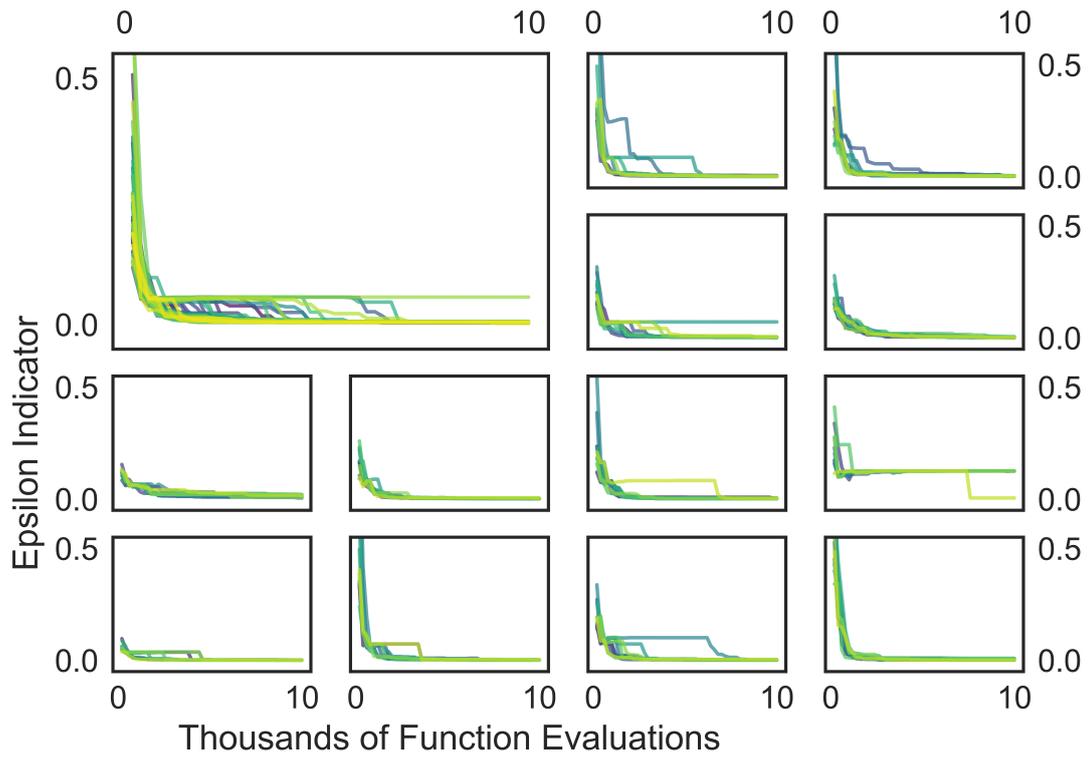


Figure S6. (top left) Epsilon indicator metric for approximate Pareto sets from Borg MOEA with 50 random seeds for baseline state of the world. (all others) Epsilon indicator metric for approximate Pareto sets from Borg MOEA with 10 random seeds each for 12 randomly selected SOWs.

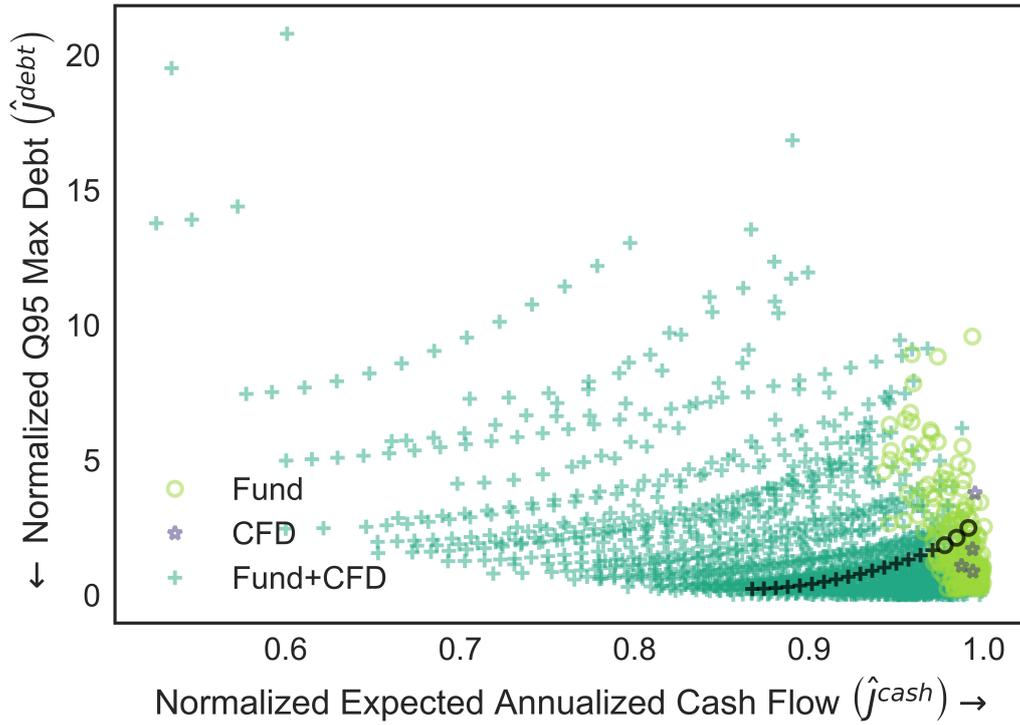


Figure S7. Same as Figure 10 from main text, but showing all unfiltered results.

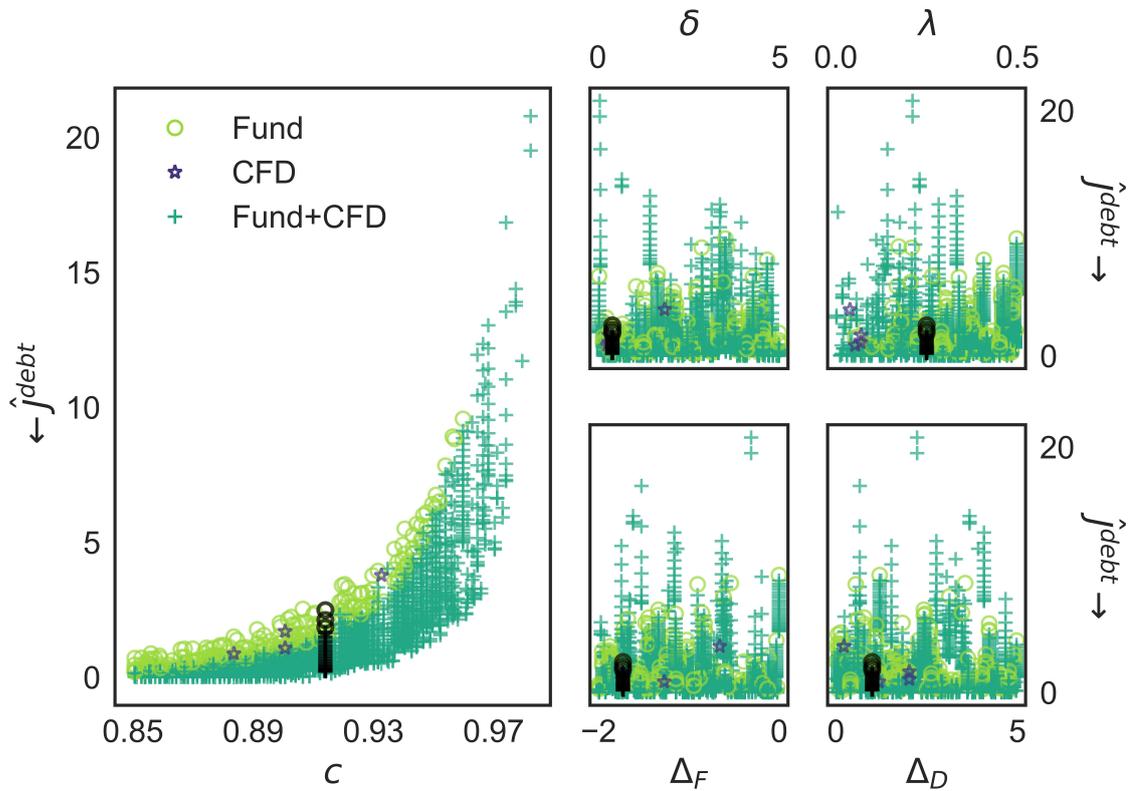


Figure S8. Same as Figure 11 from main text, but showing all unfiltered results.

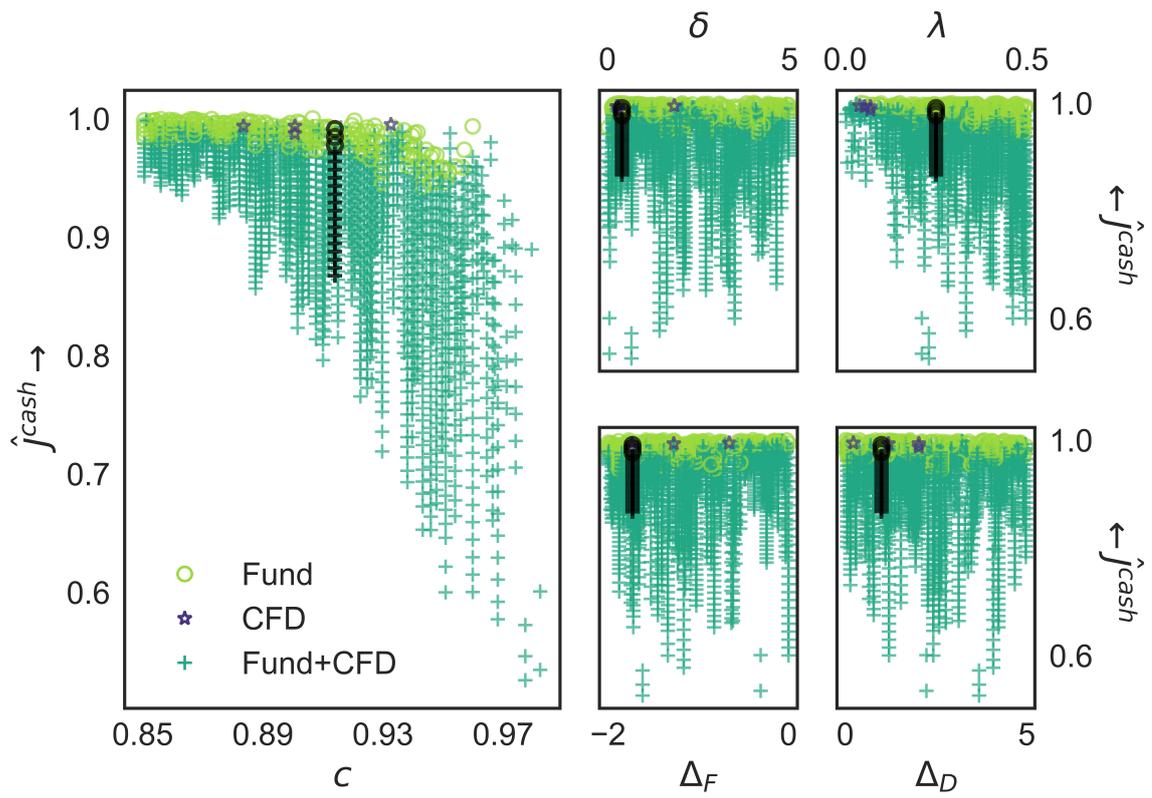


Figure S9. Same as Figure 12 from main text, but showing all unfiltered results.