Supplementary Materials for *Resource scarcity compromises explore-exploit decision-making*

Shou-An A. Chang, Julian Jara-Ettinger[1], Arielle Baskin-Sommers[1]

Yale University, Department of Psychology, New Haven, CT, USA

[1]Shared last author

Correspondence concerning this article should be addressed to Shou-An Ariel Chang, Yale University, P.O. Box 208205, New Haven, CT 06520. Email: ariel.chang@yale.edu.

Supplementary Methods and Results

Bayesian Updating of Depletion Rate

Because depletion rates consist of values between 0 and 1, we represented participant beliefs about depletion rates through a Beta distribution. Beta distributions can express a wide range of probability distributions over the range [0,1], and are determined by two parameters: α and β. This enabled us to represent beliefs about depletion rates in terms of these two parameters. Formally, upon learning a tree's depletion rate, k (obtained whenever participants harvested the same tree more than once), our model updates its beliefs about depletion rates via Bayesian inference:

$$p(\alpha,\beta \mid k) \propto p(k \mid \alpha,\beta) p(\alpha,\beta)$$

Here, $p(\alpha,\beta \mid k)$ is the updated belief that the depletion rate distribution is captured by parameters α and β (i.e., the posterior probability). This belief is given by the product of the model's prior probability over these two parameters ($p(\alpha,\beta)$), times the likelihood of finding a tree with depletion rate k under those parameters ($p(k \mid \alpha,\beta)$).

To implement this model, we used a discrete hypothesis space consisting of 400 uniformly-spaced samples, generated from different combinations of α and β parameters. The combinations of parameters included the range 1 to 20 of α values with steps of size 1, and the range 1 to 20 of β values with steps of size 1. At the beginning of the task, the prior over depletion rates were set to $p(\alpha,\beta)=1/400$ for all parameter combinations (i.e., a uniform distribution), and the likelihood function $p(k \mid \alpha,\beta)$ was given by the Beta distribution $\mathrm{Beta}(r;\alpha,\beta)$.

Finally, the predicted return rate of a new tree is given by the return rate $\overline{r}$:

$$\overline{r} = \sum_{\alpha,\beta} p(\alpha, \beta) \left(\frac{\alpha}{\alpha + \beta} \right)$$

where $p(\alpha,\beta)$ is the model's latest beliefs that the depletion rates are captured by parameters α and β, and $\alpha/(\alpha+\beta)$ is the expected return rate under those parameter combinations.

Assumption about observers' knowledge of depletion rate.

As described in our main text, our model assumes that participants will know a tree's depletion rate *k* after making consecutive harvests. However, because participants always received an integer number of apples, participant's estimated decay rates may be different from the true decay rates. To test the extent to which this assumption could bias our model, we calculated the apparent decay rate given a set of apples and compared it to the tree underlying decay rate. This analysis revealed an average mismatch of .07 and a mode mismatch of .03 (obtained by applying kernel density estimation to the list of errors, and finding the model of this distribution).

Supplemental Analyses on Patch Residence Times

A common metric used in foraging paradigms is patch residence times. In our experiment, we can use the number of "stay" decisions at a tree as a proxy for patch residence times, since each tree in our task equates to one patch, and the harvest time in our task is fixed to 3 seconds. To approximate resource-maximizing patch residence times as defined by our model in the main text, we created a fake participant who repeatedly made 6 "stay" decisions before "switching" in the traditional foraging task and who repeatedly made 4 "stay" decisions before "switching" in the depleted foraging task. Comparing this fake participant's decisions against the model predictions for the traditional foraging task and the depleted foraging task allowed us to identify for each tree how many "stay" decisions were in line with our Bayesian-based MVT model (i.e., results in a model prediction that is >= .5) and therefore how many "stay" decisions were resource-maximizing at each tree. In order to investigate each subject's patch residence times relative to MVT, for each tree, we converted “stay” decisions to time by multiplying number of “stay” decision by 3 (i.e., harvest time) and subtracted this time from the resource-maximizing patch residence time as defined by MVT. Positive values indicated a tendency for participants to underharvest relative to MVT (i.e., switch earlier than is resource-maximizing as indicated by MVT), and negative values indicate a tendency for participants to overharvest relative to MVT (i.e., switch later than is resource-maximizing as indicated by MVT).

Individual Differences in Patch Residence Times

In the traditional foraging task, the mean of the difference score between participant patch residence and patch residence time as defined by MVT was 2.05 s (*SD*=7.53), and in the depleted foraging task it was 1.95 s (*SD*=4.98), indicating that in both environments, participants typically switched to explore a new tree earlier than was resource-maximizing as defined by our Bayesian-based MVT model. Correlations between this difference score and ADI showed no significant relationship between the difference score and ADI in the traditional foraging task $r(110)=0.03$, $p = .742$ or the depleted foraging task $r(110)= -0.01$, $p = .958$, which replicates our results showing no significant relationship between our bias parameter and ADI in either foraging task.

Patch Residence Times in the Traditional and Depleted Foraging Tasks

To compare patch residence times across tasks, we ran a paired samples t-test to compare each participant’s average patch residence times in the traditional foraging task in relation to average patch residence times in the depleted foraging task. Results showed a significant difference in average patch residence times across tasks, $t(111) = 3.44$, $p < .001$, CI [2.68, 4.21], such that participant’s spent more time harvesting in patches in the traditional foraging task (*M*=12.36, *SD*=8.91) compared to the depleted foraging task (*M*=7.26, *SD*=4.81). In contrast to other foraging studies that manipulate travel time between patches (Constantino et al., 2017; Constantino & Daw, 2015; Lenow et al., 2017), we chose to keep travel time constant and manipulate patch richness across tasks in order to mimic experiences of resource scarcity. Thus, in line with research showing that individuals adjust their decision-making based on patch richness (Hutchinson et al., 2008; Wolfe, 2013; Zhang et al., 2015), we see that as predicted by MVT, participants spend more time foraging in patches richer in resources compared to patches depleted in resources, when travel time is held constant.

Results of Robustness Analyses

Below we present table summaries of the robustness analyses.

Supplemental Table 1. *Summary of Regression Results for Robustness Analyses on Resource-Maximizing Decision-Making in Resource-Rich Environment*

Predictor	*b*	*t*	95% CI Lower Bound	95% CI Upper Bound
Intercept	65.64***	48.92	62.98	68.30
Response inhibition	-.89	-.67	-3.54	1.76
ADI	**-4.59*** **	-3.55	-7.16	-2.02
R^2	.10			
Intercept	64.98***	52.16	62.51	67.45
Working memory	**-2.76** **	2.64	.88	6.22
ADI	**-2.82***	-2.09	-5.49	-.15
R^2	.14			
Intercept	65.82***	39.15	62.49	69.15
Biological sex	-2.76	-1.09	-7.87	2.28
ADI	**-3.99** **	-3.06	-6.57	-1.41
R^2	.09			
Intercept	64.86***	53.08	-6.00	-1.61
Age	**-3.80*** **	-3.43	-6.00	-1.61
ADI	**-3.49** **	-3.11	-5.72	-1.27
R^2	.16			
Intercept	66.12***	40.56	62.88	69.36
Response inhibition	.48	.30	-2.76	3.72
Working memory	**2.96***	2.11	.17	5.74
Biological sex	-3.44	-1.32	-8.64	1.75
Age	-2.87	-1.84	-5.97	.22
ADI	**-3.21***	-2.49	-5.76	-.65
R^2	.20			

Note. *p < .05, **p <.01, ***p<.00. Row 1 presents results from robust linear regression with ADI and Color-Word Stroop task Inhibition/Switching vs. Word Reading contrast z-scored as predictors. Row 2 presents results from robust linear regression with digit span backwards z-scored as predictors. Row 3 presents results from robust linear regression with ADI and biological sex dummy coded 0= male, 1=female as predictors. Row 4 presents results from robust linear regression with ADI and age z-scored as predictors. Row 5 presents results from robust linear regression with Color-Word Stroop task Inhibition/Switching vs. Word Reading contrast score, digit span backwards, biological sex, age, and ADI as predictors.

Supplemental Table 2. *Summary of Regression Results for Robustness Analyses on Resource-Maximizing Decision-Making in Resource-Depleted Environment*

Predictor	*b*	*t*	95% CI Lower Bound	95% CI Upper Bound
Intercept	63.65***	58.92	61.51	65.80
Response inhibition	-.53	.44	-2.93	1.87
ADI	**-6.08*** **	-5.49	-8.28	-3.88
R^2	.26			
Intercept	62.59***	57.01	60.41	64.77
Working memory	**2.59***	1.20	.22	4.97
ADI	**-3.83** **	-2.75	-6.59	-1.07
R^2	.19			
Intercept	62.54***	45.90	59.84	65.24
Biological sex	.25	.10	-4.69	5.19
ADI	**-4.75*** **	-3.59	-7.38	-2.16
R^2	.16			
Intercept	62.54***	59.62	60.46	64.62
Age	**-2.81** **	-2.84	-4.78	-.85
ADI	**-4.49*** **	-3.94	-6.74	-2.23
R^2	.21			
Intercept	63.02***	47.05	60.37	65.69
Response inhibition	.43	.28	-2.66	3.51
Working memory	1.86	1.59	-.47	4.19
Biological sex	.80	.35	-3.77	5.37
Age	-1.58	-1.12	-4.37	1.21
ADI	**-5.09*** **	-3.93	-7.67	-2.52
R^2	.29			

Note. *p < .05, **p <.01, ***p<.00. Row 1 presents results from robust linear regression with ADI and Color-Word Stroop task Inhibition/Switching vs. Word Reading contrast z-scored as predictors. Row 2 presents results from robust linear regression with digit span backwards z-scored as predictors. Row 3 presents results from robust linear regression with ADI and biological sex dummy coded 0= male, 1=female as predictors. Row 4 presents results from robust linear regression with ADI and age z-scored as predictors. Row5 presents results from robust linear regression with Color-Word Stroop task Inhibition/Switching vs. Word Reading contrast score, digit span backwards, biological sex, age, and ADI as predictors.

Supplemental Analyses on Resource-Maximizing Decision-Making Over Trials

In order to investigate resource-maximizing decision-making over the course of the traditional foraging task, whether or not decisions were resource-maximizing were entered in a binomial mixed-effects model with participant as a random effect and trial number z-scored as a predictor. Results showed a significant positive effect of trial number on whether or not decisions were resource-maximizing $\beta = .09$, $SE = .01$, $z(31926) = 7.29$, $p < .001$, 95% CI [.07, .11] in the traditional foraging task (see Supplemental Figure 1A). In order to investigate learning over the course of the depleted foraging task, whether or not decisions were resource-maximizing were entered in a binomial mixed-effects model with participant as a random effect and trial number z-scored as a predictor. Results showed no significant effect of trial on whether or not decisions were resource-maximizing $\beta = .03$, $SE = .02$, $z(15130) = 1.64$, $p = .101$, 95% CI [-.01, .06] (see Supplemental Figure 1B). Taken together, these results suggest that when foraging in a resource rich environment, individuals learn over time to make decisions that are resource-maximizing; however, individuals do not show learning effects when foraging in a resource depleted environment. As stated in our main paper, our model does not account for individual differences in learning, and these analyses highlight the need for future research to investigate the role of learning while individuals forage.

Supplemental Figure 1a and 1b. *Resource-Maximizing Decisions Over Trials*.

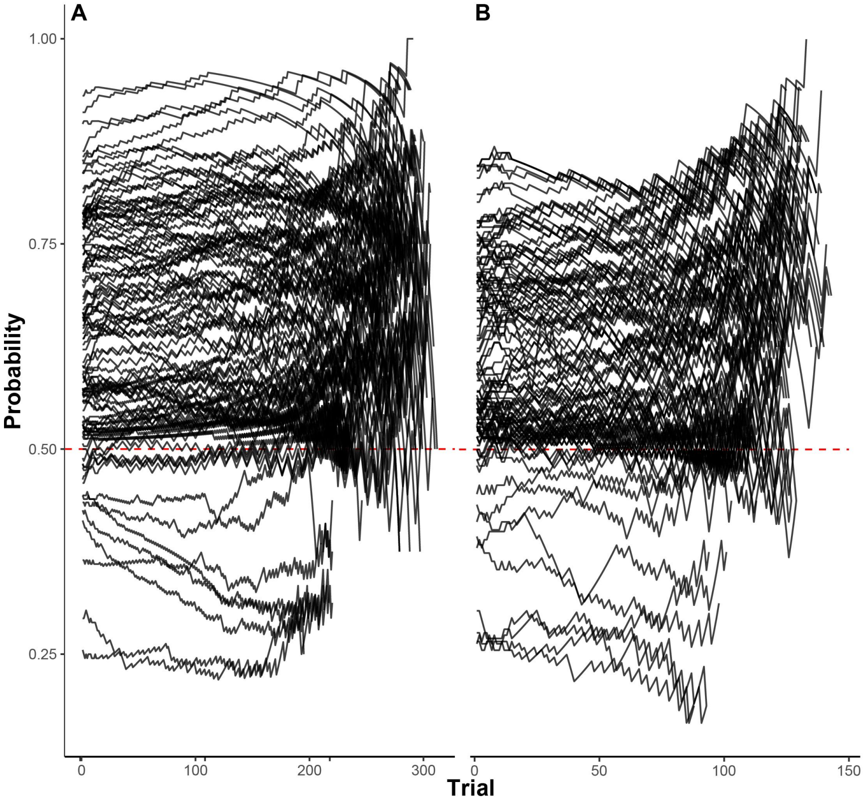

Note. Panel A shows average probabilities assigned to participant choices by the classical model over trials in the traditional foraging task. Panel B shows average probabilities assigned to participant choices by the classical model over trials in the depleted foraging task.

Supplemental Analyses for the Social Norm Foraging Task

Moral licensing theory suggests that individuals who initially behave morally may incur a "moral license" over time that increases their propensity to engage in immoral behavior at a later time point (Merritt et al., 2010). That is, individuals can draw on their past good behavior to counteract potential worries about present decisions to act immorally. In order to investigate whether or not participants experience moral licensing (i.e., are more likely to engage in social norm violations later in the task), we split the length of the social norm foraging task into four quantiles and calculated our measure of tendency to engage in social norm violation by each quantile. We entered the tendency to engage in social norm violation into a repeated measures General Linear Model with quantile as a within-subjects factor and ADI as a continuous covariate. Huynh-Feldt corrected *p*-values were reported to protect against violations of the assumption of sphericity.

Consistent with the moral licensing theory, results showed a significant main effect of quantile on tendency to engage in social norm violations $F(3, 282) = 47.33, p < .001, \eta_p^2 = 0.335$, such that participants made more decisions to cross and harvest on Logan's land later in the task (quantile 2 [M = 117.370, SD=75.35]; quantile 3 [M=125.80, SD=75.59]; quantile 4 [M=129.69, SD=76.41]) compared to the beginning of the task (quantile 1 [M=50.28, SD=31.19]). There was no significant interaction between ADI and quantile $F(27, 282) = 1.48, p = 0.062, \eta_p^2 = 0.124$ on tendency to engage in social norm violations. Taken together, and consistent with the moral licensing theory, these results suggest that participants engage in relatively limited social norm violations at the beginning of the task, which may provide participants with a "moral license" to increase engagement in social norm violations later in the task.

Alternative Computational Framework

Computational Framework. In our task, participants saw an overview of the number of apples on the trees on their land, which reflected the mean of the distribution from which initial tree richness was initialized, and they knew the harvest and travel times. It is, therefore, possible that participants in our task used these values to predict the value of exploration, rather than relying on estimates of their past return rate.

To ensure that this possibility would not affect our results, we developed a second computational model that was similar to our second model with the difference that, instead of comparing $r_{exploit} = v_{exploit} / h$ against the long-run average return rate, we compared it against

$r_{exploit} = v_{exploit} / (h + \mathrm{d})$, where $v_{explore}$ is the expected number of apples obtained by harvesting a new tree, h is the harvest time, and d is the travel time. Because participants knew all of these parameters, we set $v_{explore} = 10$ (average initial number of apples given aerial view; $v_{explore}$ was set to 6 in the depleted foraging task), $h=3$ (known harvest time), and $d=6$ (known travel time). $r_{exploit}$ was calculated in the same way as our main model.

Experimental Task Key Variables. Percentage of explore or exploit decisions the participant made that were adherent to the resource-maximizing predictions, bias, and imprecision during traditional and depleted environments were calculated using the same approach as noted in the main manuscript. Analyses are presented on 112 subjects unless otherwise specified. For analyses on the resource-decision-making parameters, 110 participants were included in the traditional foraging task and 110 participants were included in the depleted foraging task.

Alternative Model Results

Traditional Foraging Task

To examine the effect of real-world levels of concentrated disadvantage on foraging behavior, percentage of decisions that were resource-maximizing according to the alternative computational framework in a generally resource-rich environment were entered in a robust linear regression with ADI as a continuous predictor. A robust linear regression was used because after a linear regression was conducted, diagnostic plots of Cook's distance identified 7 influential points. Consistent with the model from the main text, results showed a significant effect of ADI on percentage of resource-maximizing decisions $R^2 = .07$, $b = -1.41$, $t(110) = -3.00$, $p = .003$, 95% CI [-2.35, -.48], suggesting that individuals who experience higher levels of concentrated disadvantage make fewer resource-maximizing decisions when foraging in a generally resource-rich environment.

Depleted Foraging Task

To examine the effect of real-world experience of concentrated disadvantage on foraging behavior, percentage of decisions that were resource-maximizing according to the alternative computational framework in a generally resource-depleted environment were entered in a robust linear regression with ADI as a continuous predictor. Robust regression was used because Cook's distance measures of percentage of decisions that were resource-maximizing in a generally resource-depleted environment entered in a linear regression with ADI as a continuous predictor identified 4 influential points. Consistent with the model from the main text, results showed a significant effect of ADI on percentage of resource-maximizing decisions $R^2 = .09$, $b = -1.49$, $t(110) = -3.49$, $p < .001$, 95% CI [-2.33, -.64], suggesting that individuals who experience higher levels of concentrated disadvantage make fewer resource-maximizing decisions when foraging in a generally resource-depleted environment.

Resource Decision-Making Parameters

Bias. Consistent with the model in the main text, results showed no relationship between ADI and bias parameters in the resource-rich environment, $r(101) = .04$, $p = .69$, 95% CI [-.15, .23] or in the resource-depleted environment, $r(94) = .18$, $p = .08$, 95% CI [-.02, .37].

Imprecision. Consistent with the model in the main text, results showed a significant positive relationship between ADI and the imprecision parameter in the resource-rich environment, $r(101) = .21$, $p = .035$, 95% CI [.02, .39] and a significant positive relationship between ADI and the imprecision parameter in the resource-depleted environment, $r(94) = .21$, $p = .041$, 95% CI [.01, .39].

References

Constantino, S., Dalrymple, J., Gilbert, R. W., Varanese, S., Di Rocco, A., & Daw, N. D. (2017). A Neural Mechanism for the Opportunity Cost of Time. *bioRxiv*, 173443. https://doi.org/10.1101/173443

Constantino, S., & Daw, N. D. (2015). Learning the opportunity cost of time in a patch-foraging task. *Cognitive, affective & behavioral neuroscience, 15*(4), 837-853. https://doi.org/10.3758/s13415-015-0350-y

Hutchinson, J. M. C., Wilke, A., & Todd, P. M. (2008). Patch leaving in humans: can a generalist adapt its rules to dispersal of items across patches? *Animal Behaviour, 75*(4), 1331-1349. https://doi.org/https://doi.org/10.1016/j.anbehav.2007.09.006

Lenow, J., Constantino, S., Daw, N. D., & Phelps, E. A. (2017). Chronic and acute stress promote overexploitation in decision making. *The Journal of Neuroscience, 37*(23), 5671-5689.

Merritt, A. C., Effron, D. A., & Monin, B. (2010). Moral self-licensing: When being good frees us to be bad. *Social and Personality Psychology Compass, 4*(5), 344-357. https://doi.org/10.1111/j.1751-9004.2010.00263.x

Wolfe, J. M. (2013). When is it time to move to the next raspberry bush? Foraging rules in human visual search. *Journal of Vision, 13*(3), 10. https://doi.org/10.1167/13.3.10

Zhang, J., Gong, X., Fougnie, D., & Wolfe, J. M. (2015). Using the past to anticipate the future in human foraging behavior. *Vision Research, 111*, 66-74. https://doi.org/https://doi.org/10.1016/j.visres.2015.04.003