Appendix A: Instructions

You are about to participate in an experiment in the economics of individual decision making. If you follow these instructions carefully and make good decisions, you will earn money that will be paid to you in cash at the end of the session. If you have a question at any time, please raise your hand and the research investigator will answer it. We ask that you not talk with one another for the duration of the experiment.

Task Description

You will be playing the role of a procurement manager of a firm.Your task is to procure the required 100 wodgets from the firm's suppliers. Specifically, the firm has 3 potential suppliers and you need to decide how to allocate the total procurement of 100 wodgets among these suppliers. For example, if you decide to procure from a single supplier, then you would order all 100 units from that supplier. Alternatively, if you decide to procure from 2 suppliers, then you would order 50 units from each of them.Similarly, if you decide to procure from all 3 suppliers, then you would order 33.33 units from each of them.

Supplier Disruptions

Every supplier is at risk of being disrupted by an external event (e.g. a fire in the supplier's factory). If disrupted, the supplier would not deliver any wodgets. For example, imagine you procure the 100 wodgets from one supplier, if this supplier experiences a disruption, then you will be 100 wodgets short of what you need. Alternatively, imagine you procure the 100 wodgets from two suppliers, 50 wodgets from each supplier, if one of these suppliers experiences a disruption, then in that round you will be 50 wodgets short of what you need. Similarly, imagine you procure the 100 wodgets from three suppliers, $33.\overline{3}$ wodgets from each supplier, if one of these suppliers experiences a disruption, then in that round you will be $33.\overline{3}$ wodget short of what you need.

You cannot influence whether a supplier is disrupted or not. But you do know that any supplier is disrupted with a probability of 1%.

Total Procurement Cost

Your supplier selection affects two cost components. First, you incur a fixed cost of $\$33.\overline{33}$ for every supplier that you procure wodgets from in a given round:

$$\text{Total Fixed Cost} = \$33.\overline{33} * \text{Number of suppliers you procure from}$$

If you procure from one supplier, then you incur a total fixed cost of $\$33.\overline{33}$. If you procure from two suppliers, then you incur a total fixed cost of $\$66.\overline{66} = \$33.\overline{33} * 2$. If you procure from three suppliers, then you incur a total fixed cost of $\$100 = \$33.\overline{33} * 3$. Importantly, you incur the fixed cost even if the supplier is disrupted and does not deliver any wodgets.

Second, you incur a **penalty cost** when any of your selected suppliers is disrupted and does not deliver wodgets. You incur a penalty cost for every wodget short of the 100 wodgets that you need, for not being able to produce the final product for your customers. The more wodgets you fail to procure because of disruptions at your supplier(s), the larger the total penalty cost. The total penalty cost increase is quadratic in the number of wodgets not received:

$$\text{Total Penalty Cost} = (\$1.\overline{01} * \text{Number of wodgets not received})^2$$

For example, imagine you selected 3 suppliers, and one is disrupted. Then you are short $33.\overline{3}$ wodgets, and you incur a total penalty cost of \$1122.33 ($= \$1.\overline{01} * 33.\overline{3}^2$). Alternatively, imagine you selected 2 suppliers, and one is disrupted. Then you are short 50 wodgets, and you incur a total penalty cost of \$2525.25 ($= \$1.\overline{01} * 50^2$). Similarly, imagine that you selected 1 supplier that is disrupted (or that you selected 3 suppliers all of which are disrupted), then you are short 100 wodgets, and you incur a total penalty cost of \$10,101.01 ($= \$1.\overline{01} * 100^2$).

Your goal in this experiment is to minimize the total procurement cost, which is the sum of the fixed cost and any penalty cost incurred in each round:

$$\text{Total Procurement Cost} = \text{Total Fixed Cost} + \text{Total Penalty Cost}$$

Number of rounds

After you have made your supplier selection in a round, then you move on to the next round, where you will source from a new set of suppliers.It is important to note that the suppliers in each round are unique and a supply disruption (or lack thereof) in a given round has no influence on the outcome of any other rounds. You will play a total of 200 independent rounds.

How we determine your payment

At the end of the experiment, the computer will calculate your **total earnings**, by deducting the total average procurement cost that you have **accumulated across the 200 rounds**, from an endowment of \$365 laboratory dollars given to you at the beginning of the experiment. The total earnings will then be converted to US dollars. Specifically, you will be paid \$1.00 US dollars for every \$16.83 laboratory of your total earnings in the experiment. On the final screen you will be able to see your total earnings in US dollars for this session. You will be paid in cash at the end of the session. All earnings are confidential, though you will have to sign a sheet indicating how much you have been paid.

Appendix B: Interface

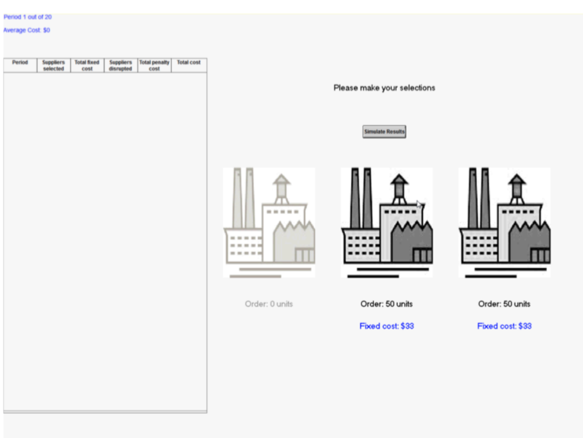

Figure 8: User Interface: Sourcing Decision

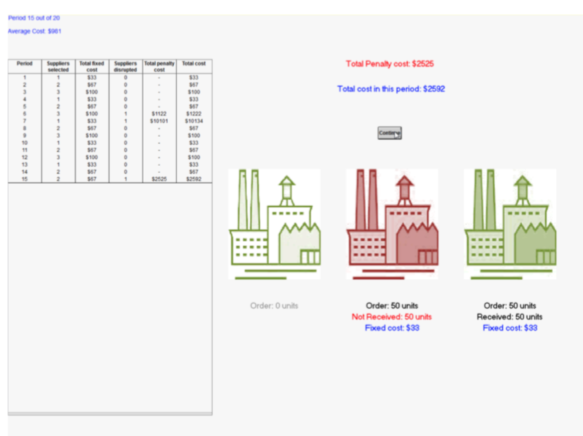

Figure 9: User Interface: Subject Feedback

Appendix C: Proofs and Derivations

C.1. Derivation of the expected cost function (Equation 2)

Let $k(n)$ be the number of failed deliveries. With D/n units per delivery, the number of units short is $k(n)(D/n)$, and the expected value of the cost of missed sales defined in Equation (1) is

$$\mathbb{E}L\left(\frac{k(n)D}{n}\right) = \left(\frac{\alpha D^2}{n^2}\right)\mathbb{E}[k(n)^2].$$

With each of n suppliers having independent probability of failed delivery p, the number of failed deliveries, $k(n)$ is binomially distributed with mean np and variance $np(1-p)$, so $\mathbb{E}[k(n)^2] = np(1-p) + (np)^2$, and

$$\begin{aligned}\mathbb{E}L\left(\frac{k(n)D}{n}\right) &= \left(\frac{\alpha D^2}{n}\right)[np(1-p) + (np)^2] \\ &= \left(\frac{\alpha D^2}{n}\right)[p + (n-1)p^2]\end{aligned}$$

which is the second term in the expected cost function, Equation (2).

C.2. Proof of Proposition 1

The first part of the proposition states that expected cost for any fixed choice of n, with δ held constant, is increasing in ω and p. As noted in Section 3.1, $C(n)$ is defined only on positive integers. We restate the expected cost function presented in Equation (2)

$$C(n) = \delta\left[n + \left(\frac{\omega}{n}\right)[p + (n-1)p^2]\right].$$

Without loss of generality, we can set $\delta = 1$ and focus on the terms inside the brackets. It is straightforward to show for any fixed value of n, the cost function is increasing in p, ω.

$$\begin{aligned}\frac{\partial C(n)}{\partial p} &= \left(\frac{\omega}{n}\right)[1 + 2(n-1)p] > 0 \\ \frac{\partial C(n)}{\partial \omega} &= \left(\frac{1}{n}\right)[p + (n-1)p^2] > 0\end{aligned}$$

We examine how the rate of cost increase with these parameters is affected by n:

$$\begin{aligned}\frac{\partial^2 C(n)}{\partial n \partial \omega} &= -\frac{\omega p(1-p)}{n^2} \leq 0, \\ \frac{\partial^2 C(n)}{\partial n \partial p} &= -\frac{\omega}{n^2}[1-2p] \leq 0 \text{ for } p \leq 0.5.\end{aligned}$$

Since the above second partial derivatives are negative, the *rate* of cost increase is lower for higher n, meaning an increase in these parameters cannot make a smaller value of n more attractive. The intuition for the constraint $p \leq 0.5$ is that when the probability of disruption is high, it is not attractive to incur additional fixed cost to diversify with suppliers who are not very reliable. As we are interested in low-probability events, we restrict our study to cases where $p \leq 0.5$.

□

C.3. Proof of Proposition 2

For any given n, define $\omega(p)$ as a function of p where $\omega(p)$ will change as needed to maintain the symmetric property $C(n-1)=C(n+1)$. By convexity of $C(n)$, $C(n-1)=C(n+1)$ implies that n is optimal. Setting the difference in cost between $C(n-1)=C(n+1)$ to zero and rearranging yeilds:

$$C(n+1)-C(n-1) = 2\delta\left[1-\frac{\omega(p)p(1-p)}{(n^2-1)}\right]=0$$
$$\omega(p) = \frac{(n^2-1)}{p(1-p)}.$$

□

C.4. 2 supplier design details (Section 4.4.1)

We select values for $C(1), C(2)$, and p. These three values uniquely determine α and δ.

$$\begin{aligned} C(1) &= \delta+\alpha D^2 p \\ \delta &= C(1)-\alpha D^2 p \\ C(2) &= 2\delta+\alpha D^2 p\left(\frac{1+p}{2}\right) \\ &= 2\left(C(1)-\alpha D^2 p\right)+\alpha D^2 p\left(\frac{1+p}{2}\right) \\ &= 2C(1)-\frac{\alpha D^2 p(3-p)}{2} \end{aligned}$$

Appendix D: I-SAW details (Section 5)

Probability of inertia is given as

$$\Pr(\text{Inertia}_{t+1}|\text{not in exploration mode}) = \pi_i^{Surprise_t}.$$

Let $C_t(n)$ denote the realized cost in period t given a choice of n suppliers, and define

$$\bar{C}_t(n) = \frac{1}{t}\sum_{i=1}^{t} C_i(n).$$

to be the grand mean of all past outcomes. The expression to measure deviations from current and past outcomes equally weights (a) gaps between t and $t-1$ outcomes, and (b) gaps between outcomes in t and the grand mean of all past outcomes,

$$Gap(t) = \frac{1}{6}\left[\sum_{n=1}^{3}|C_t(n)-C_{t-1}(n)| + \sum_{n=1}^{3}|C_t(n)-\bar{C}_{t-1}(n)|\right]. \tag{1}$$

To initialize, we set $t=0$ cost outcomes to be equal to expected costs. Similarly define the grand mean of past values of *Gap* as

$$\bar{G}_t = \frac{1}{t}\sum_{i=1}^{t} Gap(t)).$$

We normalize $Gap(t)$ to guarantee that it is between 0 and 1 by dividing by the sum of the current gap and the average of past gap values. Following Nevo and Erev (2012) we call this normalized term *Surprise*.

$$Surprise_t = \frac{Gap(t)}{Gap(t)+\bar{G}_{t-1}} \tag{2}$$

D.1. I-SAW estimation details

While Nevo and Erev (2012) use a simulation and grid search method to find parameter values of their model, we adopt a sample average approximation (SAA) approach where a pool of 1,000 (virtual) subjects are created using latin hybercube sampling, and the parameters are optimized on this sample. This approach, particularly with latin hypercube sampling, is efficient in solving stochastic optimization problems (Freimer et al., 2012). As is typical with SAA, fit is then evaluated on a new sample of subjects (10,000 in our case) to create an unbiased fit estimate.

D.2. Alternative I-SAW models

We estimate I-SAW parameters for each of the four conditions separately. First, as it must, SSE is lower when we fit condition by condition as compared to the pooled model. Total SSE drops from 5.524 (Table 5) to 5.042 (Table 7), a 9.6% reduction. We interpret these results with caution for several reasons. First, this fit improvement comes with a 4X increase in parameters. Second, much of the improvement comes from the better fit for condition *High* which is characterized by many repeated single-sourcing decisions which can be captured by either inertia or small-sample-biased exploitation decisions. Also, this is the only condition where the observed frequency of triple sourcing is nearly zero for many periods.

Cond.	ϵ_u	π_u	w_u	ρ_u	μ
pooled	0.760	0.942	0.068	0.627	1
High	0.521	0.922	0.249	0.771	1
Low	0.390	0.886	0.603	0.958	4
$High^p$	0.434	0.921	0.140	0.849	3
$High^\omega$	0.746	0.949	0.063	0.858	4

Table 6: I-SAW Learning Model: Pooled and Condition Estimation Results

	High $(p=1\%)$ $n^*=2, \bar{n}=1.35$			*Low* $(p=20\%)$ $n^*=2, \bar{n}=1.71$			$High^p$ $(p=5\%)$ $n^*=3, \bar{n}=1.63$			$High^\omega$ $(p=1\%)$ $n^*=3, \bar{n}=1.56$		
Mode	Explore	Inertia	Exploit	Explore	Inertia	Exploit	Explore	Inertia	Exploit	Explore	Inertia	Exploit
Proportion	26%	48%	25%	20%	45%	36%	22%	48%	30%	38%	41%	21%
$\hat{n}$	1.57	1.36	1.10	1.91	1.74	1.55	1.95	1.66	1.34	1.83	1.55	1.11
MAD		0.039			0.040			0.033			0.032	
SSE		1.356			1.575			1.120			0.991	

Table 7: I-SAW Learning Model: Average sourcing by decision mode (using condition parameter estimates)

Next, starting from our pooled model fit, we fit I-SAW with one of the three decision modes removed. As shown in Table 8, the presence of each of the decision modes is helpful to the overall fit. With no exploration, subjects are either making an exploitation decision or repeating that decision while in inertia mode. Removing inertia is the least damaging to fit since repeat, consolidation decisions can to some extent be captured by exploitation mode. With no exploitation mode, all decisions follow round 1 decisions since decisions are either exploration decisions or repeat of exploration decisions.

Model	ϵ_u	π_u	w_u	ρ_u	μ	SSE (Pooled)
Baseline	0.760	0.942	0.068	0.627	1	5.525
No Exploration	n/a	0.994	0.521	0.176	1	35.259
No Inertia	1.00	n/a	0.125	0.078	1	7.218
No Exploitation						129.302

Table 8: I-SAW Learning Model: Pooled Estimation Results Removing Decision Parameters

D.3. Additional I-SAW counterfactual studies

As indicated in Figure 7, an increase in μ, the parameter that governs the size of a small sample that a decision maker considers, mitigates consolidation bias in exploitation decisions, with this mitigation substantially more effective for lower severity settings. Figure 10 tells a similar story regarding increasing the weight decision makers place on the entire history when making exploitation decisions. It is worth recalling that (1) these simulations are for the same 200 periods from our experimental studies, and (2) w_i for a particular subject is drawn from $(0, w)$. So, even with $w = 1$ decision makers will not in general weigh the entire history.

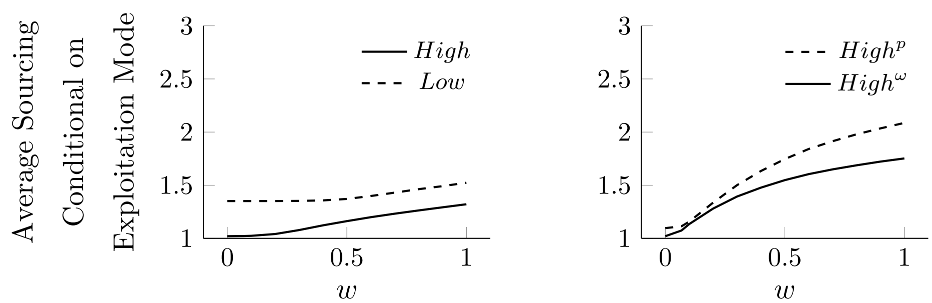

(a) Study 1 (b) Study 2

Figure 10: Effect of increasing w (weight on entire history) on average sourcing decisions *conditional on being in exploitation mode* over the same 200 rounds from our experimental studies. Increasing w reduces consolidation bias, particularly for less-severe conditions, but decisions are still well short of optimal.

Figure 7 shows that increasing the small sample parameter μ mitigates consolidation bias. In Figure 11 we see the small, negative effect of recency bias and how that bias interacts with the value of μ. In Figure 11 we repeat the two pairs of lines from Figure 7 and add comparable lines where recency bias was removed by eliminating ρ and drawing samples for a decision maker with all prior periods equally likely. As in the main I-SAW model, this sampling is still done with replacement. Figure 11 indicates that recency bias mildly exacerbates consolidation bias. In our main pooled model, the recency parameter is $\rho = 0.63$, and a subject's recency bias (ρ_i, the likelihood they select period $t-1$) is drawn from $(0, \rho)$. Recency bias exacerbates consolidation since a decision maker with recency bias drawing with replacement will over-sample period $t-1$, and a decision maker without recency bias will get a slightly more representative sample (more independent observations) for the same small sample size. As indicated in the Figure, this effect is relatively small in our settings. So, while finding interventions to mitigate recency bias in decision makers is helpful, it may not be the top priority.

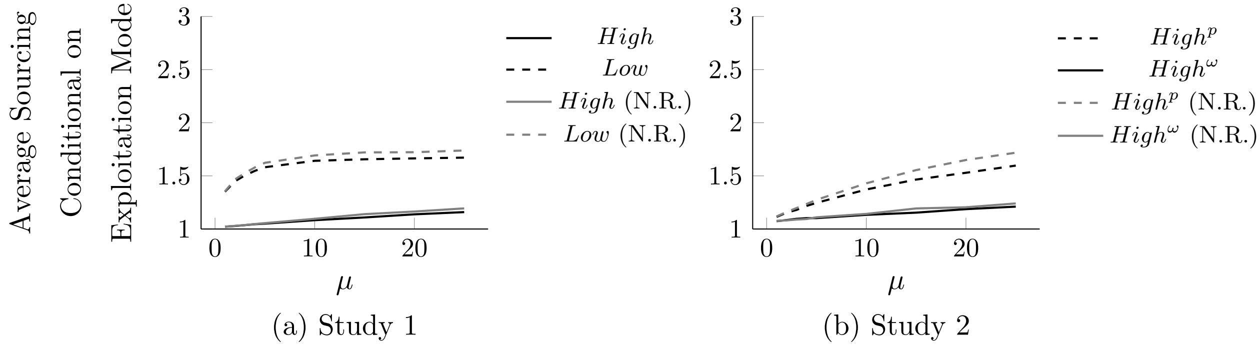

Figure 11: Effect of μ (small-sample bias) and recency bias on average sourcing decisions *conditional on being in exploitation mode* over the same 200 rounds from our experimental studies. Increasing μ reduces consolidation bias. Recency bias exacerbates consolidation bias due to decision makers over-sampling the $t-1$ period.

References

Freimer, Michael B, Jeffrey T Linderoth, Douglas J Thomas. 2012. The impact of sampling methods on bias and variance in stochastic linear programs. *Computational Optimization and Applications* **51**(1) 51–75.

Nevo, Iris, Ido Erev. 2012. On surprise, change, and the effect of recent outcomes. *Frontiers in Psychology* **3** 24.