Appendix A: Model details[1]

This section first sketches how the model of CSW's environment operates in each period and the actions taken by each agent. Then each model segment (4 loops over the set of agents per period) is explained in detail to give the curious reader comprehensive insight into how the agents learn (or fail to learn) specialization and exchange with an appropriately specialized partner. Unless otherwise noted, all parameters are drawn directly from CSW.

A.1 Model overview

The pseudocode displayed below gives a general view of how each period transpires. The four loops of each period will be explained in more detail below. The data collected will be used to compare various parameterizations of the model to one another and to gauge the extent to which the model accurately approximates human behavior.

Model Pseudocode:

Set Global Parameters
Initialize Agents
Begin Loop Over Periods
    Begin Period
        Loop 1—Production and Trading Partner Selection
        Loop 2—Autarkic Consumption
        Loop 3—Trade and Update Trade Probability
        Loop 4—Update Learning, Specialization, and Willingness to Trade
    End Period
    Record Data
End Loop Over Periods

Each simulation contains $N = 8$ agents and lasts for 35 periods of specialization, exchange and consumption. Agents are divided into two types, distinguished by their respective preferences and production functions. In every period, agents produce goods (red and blue) according to their production functions and their adaptive choice of specialization rate. Then, each agent chooses at random (from a private, incrementally-updating probability distribution) one other agent as a prospective trading partner. Agents privately consume their stocks to maximize autarkic earnings and store the remaining goods for the prospect of exchange.[2]

If two agents agree to exchange (that is, if two agents each initiate a trading relationship with the other), they trade their goods to minimize total waste and then update their trading-partner-selection probability distributions on the basis of the

[1] This appendix is copied from Sect. III of Kimbrough (forthcoming).

[2] While in CSW, all goods are *technically* consumed at the end of a period (and thus, after exchange), this choice reflects the observation that incompletely specialized subjects generally trade only marginal, leftover units. Note: when agents are fully specialized in one good, nothing is consumed here and all goods remain for use in exchange.

outcome. Then, agents observe their total earnings and employ this information to "learn" whether previous incremental adjustments of specialization were effective. They update their rates of specialization for the next period, adjust the parameters of their learning rules, and the period is over. In contrast to a model of perfectly rational agents, in which agents weigh the relative payoffs of completely specified alternatives, these software agents learn specialization and exchange by employing a number of learning heuristics meant to capture the incremental process by which humans explore and master their environment.

These learning rules suggest a relatively simplistic, boundedly rational trial and error process of exchange and specialization, and they build in the possibility of "giving up" on trade and adopting an autarkic stance if exchange attempts prove fruitless. Agents compare their post-exchange earnings to parameters meant to describe their risk preferences and discount rate. The origin of these parameters is not accounted for in the model, but they could be described in the literature variously as resulting from concerns for reciprocity, fairness, or social preferences. The incremental nature of change and the fact that decision rules are tied to other *individuals* suggests a reciprocity interpretation. My agents are not concerned about the payoffs of other agents.

A.2 Model setup

At the beginning of each simulation, global parameters are instantiated to define the model. Then, a number of parameters must be assigned to each agent. Some of these parameters help agents make their decisions and may be altered via interactions with other agents and with the environment. The pseudocode below lists these relevant variables and how they are initialized. I will refer back to these variables in explaining the model below and will italicize them to ease the reader's recognition.

Setup Pseudocode:
Set Global Parameters
    #_OF_PERIODS ($T = 35$)
    #_OF_AGENTS ($N = 8$)
    Specialization Increment ($s = 0.6$)
    Trade Probability Increment ($v \in \{0.1, 0.2, \mathbf{0.3}, 0.4, 0.5\}$)[3]
    Conservatism Increment ($x = v/2$)
    Minimum Earnings in Autarky ($e^T \in \{15, 20\}$)
    Minimum Earnings in Trade ($e^A \in \{10, 15, 20\}$)

Initialize Agents
Begin Loop Over Agents ($i \in I$)
    Set ID $= i$
    Set Type and Preferences

[3]Grid search led me to a default increment of $v = 0.3$, which (at granularity of 0.1) is the weakest level of reinforcement for which simulated data is indistinguishable from human subject data. At 0.2 or 0.1, agents do not sufficiently focus on (or trust) a single trading partner to develop effective trading relationships. The results reported below are robust to values between 0.3 and 0.5.

Agents have integer constrained Leontief preferences
(U_{Type} over *red* (r) and *blue* (b))
If i mod $2 = 0$
Type = *even*
$U_{even} = \min\{2r, b\}$
Else
Type = *odd*
$U_{odd} = \min\{r, 3b\}$
End If
Set Endowment ($E_i = 0$)
Set Earnings ($e_{i,t} = 0$)
Set Specialization Rate ($\tau_i \in [0, 10]$)[4]
$\tau_i = 5$
Set Learning Rule ($l_i =$ *far-sighted* or *short-sighted*)
$l_i \sim B(1, 0.5)$, $l = 1 \Rightarrow$ *far-sighted*
Set Direction of Specialization Rate Change ($d_i = 1$ or -1)
$d_i \sim B(1, 0.5)$
Set Trade Vector ($\Delta_i = \{\delta_{i1}, \delta_{i2}, \ldots, \delta_{ij}\}$)
$\delta_{i1} = \delta_{i2} = \cdots = \delta_{ij} = 1/(N-1)$, $\delta_{ii} = 0$, and $\sum_{\in I} \delta_{ij} = 1$
Set Willingness to Trade ($w_i =$ *willing* or *unwilling*)
$w_i \sim B(1, 0.66)$, $w = 1 \Rightarrow$ *willing*[5]
Set Conservatism ($c_i \sim U(0, 1)$)
Set Probability of Mimicking ($m_i \sim U(0, 1)$)
End If
End Loop Over Agents

A.3 Loop 1—production and trading partner selection

In the first loop agents produce goods and choose prospective trading partners. In the production phase, each agent is endowed with $T = 10$ seconds of production time and based on its *Type* and choice of specialization rate, τ, produces according to the relevant production function specified in the pseudocode below. Agents store their red and blue production for later phases of consumption and exchange.

After producing goods, each *willing* agent i selects a prospective trading partner j with probability δ_{ij} from its Trade Vector. The initial probability of choosing any given agent as an exchange partner is equal for each other agent. This suggests that agents have no prior reason to choose any one agent over another, and the probability of an agent attempting to trade with itself is zero for obvious reasons. If two agents each choose one another in the same period, then the agents will exchange later in

[4]CSW allow their subjects to set their levels of specialization after exploring their production functions for a few minutes before the experiment. If subjects do not choose their own value of τ, a 50–50 split is the default. Additional, unreported simulations display path-dependence with negative efficiency consequences when all agents begin a simulation at $\tau = 2.5$ or 7.5.

[5]2/3 are initialized as willing because 1/3 of subjects in a 2-person version of CSW never engage in trade.

the period. This method of developing exchange relations can be described as a two-sided stochastic discovery process by which agents on either side of a prospective exchange sample the other available agents as potential trading partners until a pair share similar goals and are able to initiate a trade.[6]

Trading Partner Selection and Production Pseudocode

Begin Loop 1—Production and Trading Partner Selection

Set Production (R, B)

Each type produces with increasing returns to one of the two available goods, *red* and *blue*.

If Type = odd

Set $R_{i,t} = \frac{13}{10\sqrt{10}}\tau^{\frac{5}{2}} \approx 0.41\tau^{\frac{5}{2}}$

$B_{i,t} = \frac{10}{10-(\frac{300\sqrt{10}}{13})^{\frac{2}{5}}}(10-\tau) \approx 2.25(10-\tau)$

Else

Set $R_{i,t} = \frac{13}{10-\sqrt{\frac{260}{11}}}\tau \approx 2.53\tau$

$B_{i,t} = \frac{11}{10}(10-\tau)^2$

End If

Choose Potential Trading Partner

Select ϕ_i from Δ_i with $p(\phi_i = 1) = \delta_{i1}$

If $\phi_i = j$ & $\phi_j = i$

Form Trading Pair (i, j)

End If

End Loop 1—Production and Trading Partner Selection

A.4 Loop 2—autarkic consumption

Given an agent's *Type*, agents maximize consumption over the relevant utility function. The remaining goods that cannot be consumed in appropriate proportion are set aside for trading (if possible) and may be wasted if no trade occurs (since unconsumed goods do not carry from one period to the next).[7]

Autarkic Consumption Pseudocode

Begin Loop 2—Autarkic Consumption

Set Earnings (e)

If Type = even

$e_{i,t} = \max(U_{even})$ s.t.

$r = R_{i,t}$, $b = B_{i,t}$, and $e_{i,t} \in \{z^+\}$

[6]Tesfatsion (1997) develops an agent-based "trade network game" with endogenous partner selection using a variation on the Gale and Shapley (1962) mechanism to select trading pairs who then make decisions in a prisoner's dilemma. In contrast to my model, those agents may make offers to multiple potential trading partners.

[7]I chose this method because it maps nicely into behavior observed in CSW, where subjects tend to first consume all they can of what they have produced and *only then* begin to look for opportunities to gain from exchange.

Else

$e_{i,t} = \max(U_{odd})$ s.t.

$r = R_{i,t}$, $b = B_{i,t}$, and $e_{i,t} \in \{z^+\}$

End If[8]

Store Remaining Goods for Trade

Reduce Production by the Amount Consumed

Set $R_{i,t} = R_{i,t} - r$ consumed

$B_{i,t} = B_{i,t} - b$ consumed

End Loop 2—Autarkic Consumption

A.5 Loop 3—trade and update trade probability

To carry out trades among the pairs created in Loop 1, I employ a bilateral process similar to that which generates the decentralized non-monetary general equilibrium described in Starr (1976).[9] When two agents exchange, their leftover goods are pooled, and an integer-programming problem minimizes total waste and distributes the goods across the agents' utility functions accordingly.[10] All remaining goods that are not consumed either privately or as a result of exchange are wasted and destroyed at the end of the period. Note that this rules out an agent having multiple trading partners in a single period.[11]

Trade Pseudocode

Begin Loop 3—Trade and Update Trade Probability

For each trading pair (i, j)

Get Agent *Types* $(Type_i, Type_j)$

Get Leftover Goods $(R_{i,t}, B_{i,t}, R_{j,t}, B_{j,t})$

Do Trades

Minimize Waste Given Types and Leftover Goods

If $Type_x = odd$ & $Type_y = even$

$\max(r_i + b_j)$ s.t.

[8] In autarky, *odd* (*even*) agents optimally spend 56% (51%) of their time producing red to create 30 (13) reds and 10 (26) blues. This yields earnings of 30 per period for *odd* agents and 26 per period for *even* agents.

[9] I discovered Starr's paper late in this project while trying to understand the implications of subjects' trading intuition and was struck by how well the model matched up with subjects' apparent logic. In Starr's model, each agent trades with each other agent *repeatedly* until stocks have been exhausted and equilibrium is achieved. However, an economy operating under this principle is not guaranteed to converge in finite time. Here I use only a single-exchange approximation of his process.

[10] CSW often observe subjects making statements like, "I need x blue so I don't waste any red." Furthermore, subjects often come to the idea of trade because they have leftover goods that would otherwise be wasted. Thus, waste-minimization seems a reasonable assumption about how agents approach the opportunity to exchange.

[11] CSW observe that exchange is largely (though not completely) bilateral and becomes increasingly so over time. It would be possible to extend this model to second- and third-order exchange relationships, but that route is not pursued here, though the model would then be closer to Starr (1976). However, leftover goods from the first optimization tend to be small and secondary trades would often be of little marginal value.

$$\begin{aligned}
&r_i + r_j \le R_{i,t} + R_{j,t}\\
&b_i + b_j \le B_{i,t} + B_{j,t}\\
&3b_i - r_j = 0\\
&2r_j - b_j = 0\\
&r_i, r_j, b_i, b_j \in \{z^+\}
\end{aligned}$$

Update Earnings after Trade

Set $e_{i,t} = r_i + e_{i,t}$

$e_{j,t} = b_j + e_{j,t}$

End If[12]

Then, using a rule similar to "interactive reinforcement" as described in Skyrms (2004), agents adjust their *TradeVectors* to reflect the outcome of their exchange, introducing a notion of acceptable risk and a willingness of agents to simply give up on trade altogether.

Update Trade Probabilities Pseudocode

Update Probabilities of Future Trade (Skyrms 2004)

For each trading agent i and its trading partner j

If $e_i \ge e^T$

Set $\delta_{ij} = \delta_{ij} + v$

$\forall \delta > 0 \in \Delta_i \backslash \delta_{ij}$

$\delta = \delta - v/\Delta_i \backslash \delta_{ij}$[13]

Else If $e_i < e^T$

[12] The optimization is applied similarly for different trading pair compositions. Of interest to both price theorists and game theorists, the program has, by default, perfect knowledge of agent preferences and employs a computational method that far exceeds in complexity anything that human subjects can perform mentally. The most efficient allocation of goods is computed and the goods are distributed among the traders accordingly, without regard to who possessed what before the exchange took place. Thus, strategic and computational issues that may impact real-world exchanges are simplified away. I want to be clear that my choice of price determination method avoids the subtle implication that subjects in an experiment solve the price determination problem in this way. In general, economics assumes that people are intelligent actors who seek to optimize consumption relative to their preferences. While people may not use these algorithms to solve the problem in practice, their behavior is supposed to approximate the algorithmic solution just the same. However, if the social nature of exchange matters or the relative information sets of economic agents differ, exchange prices and quantities may differ, and hypotheses on these factors are interesting points that merit study in their own right. For example, if someone hypothesizes that concerns for fairness play a role in people's exchange decisions, I could impose a rule requiring agents to refuse exchanges that disproportionately benefit the other party. Or if one believes that strategic factors impact bargaining behavior, rules could be imposed requiring agents to bargain over their exchange rate in proportion to their production or the number of trade relationships available to them. On the other hand, if the concern were that the intelligence of the agents too far exceeds that of real-world actors (in that humans can't be expected to perform rapid optimizations), one way to address this would be to add noise to the exchange price derived from the optimization. The integer constraint is drawn from CSW.

[13] If at this point, any $\delta_i \le 0$, the $|\delta_i|$ is redistributed equally over the remaining positive, non-zero probability agents. The decrementing process differentiates my reinforcement from that of Skyrms because it allows some δ_i to fall to zero, and once those probabilities reach zero, they remain so indefinitely. Agents only apply the incrementing/decrementing process to *positive probability* agents in Δ_i. Hence, if $\delta_{ij} = 1$ (meaning all other δ_i are 0) and $e_{i,t} < e^T$ then the agent will decrement δ_{ij} without increasing the probability of trading with another agent and eventually all probabilities could fall to 0. If all δ_i fall to zero, the agent becomes an autarkist, i.e. unwilling to trade and short-sighted in exploring its production function.

Set $\delta_{ij} = \delta_{ij} - v$
$\forall \delta > 0 \in \Delta_i \backslash \delta_{ij}$
$\delta = \delta + v/\Delta_i \backslash \delta_{ij}$
End If
End Loop 3—Trade and Update Trade Probability

A.6 Loop 4—parameter updates

Next, on the basis of earnings in the present period, each agent makes decisions about its method of learning, rate and direction of specialization, and willingness to trade. Throughout each session, agents track their best autarky earnings and record the associated level of specialization, τ^*. If an agent is not involved in exchange in the present period, it compares its earnings e to its previous best in autarky and updates that memory if the present period exceeds its current value.

A.6.1 Learning rule

Next, each agent evaluates the effectiveness of its learning rule, l. Depending on whether an agent has traded or acted as an autarkist, it compares its earnings to e^T or e^A and decides whether to change the rule. See the pseudocode below.

Parameter Update Pseudocode
Begin Loop 4—Parameter Updates
Update Learning Rule
If i traded in period t
If $e_{i,t} \geq e^T$
Set $l_i =$ *far-sighted*
Else If $e_{i,t} < e^T$
Set $l_i =$ *short-sighted*
$\tau_i = \tau^*$
$c_i = c_i - v/2$
If $e_{i,t-1} < e^T$, $w_i =$ *unwilling*
End If
Else If i was an autarkist in period t
If $e_{i,t} < e^A$
Set $l_i =$ *short-sighted*
$\tau_i = \tau^*$
$c_i = c_i - v/2$
End If
End If

A.6.2 Specialization

Next, each agent applies its learning rule to adjust its specialization rate, τ, via a modified hill-climbing process. The two learning rules, each of which can be thought of as describing an agent's discount rate and/or level of risk-aversion in hill climbing, guide the agents in adjusting their rate of specialization.

Parameter Update Pseudocode (continued)

Update Specialization Rate

If $l_i =$ *far-sighted*

If $e_{i,t} < e^A$, $d_i = d_i^* - 1$

Else If $l_i =$ *short-sighted*

If $e_{i,t} < e_{t-1}$, $d_i = d_i^* - 1$

End If

Set $\tau_i = \tau_i + (d_i * s)$[14]

A specific example of how short-sightedness may detrimentally impact the development of specialization and trade (even between partners who ought to be trading) can be seen in Fig. 1. The z-axis represents the profit of an *odd* agent in the simulation, given its own rate of specialization and that of an *even* agent with which it is exchanging.[15] An *odd* agent that is incrementally increasing its rate of specialization while its trading partner is already fully specialized (or nearly) may actually experience a decrease in earnings at one stage in the process relative to the previous period. Notice the valley in the back left portion of the figure. Here the algorithm takes goods from an under-specialized *odd* agent and transfers them to the more specialized *even* agent with which it is trading. Although joint profits are maximized (see Fig. 2) and steadily increasing as odd specialization increases, for levels of specialization near perfect autarky, the *odd* agent experiences a decrease in welfare relative to a case in which it is less specialized. In such cases, an odd agent with short-sighted l would reverse the direction of its increase in specialization, and might never achieve the optimal outcome.

Path-dependence resulting from the initial rates of specialization of new trading partners (initial location on the profit landscape) may prevent a given pair from converging to the global optimum, instead stranding them on a local peak or breaking their trading relationship apart before it can become mutually beneficial. An agent employing the far-sighted learning rule, on the other hand, will more likely ignore this temporary decrease and continue specializing.

A.6.3 Willingness to trade

Finally, if other agents in the simulation are trading in this period, unwilling or short-sighted agents may decide reopen the possibility of trading and also to copy the production decisions of more successful agents. They learn from successful examples to

[14] Note that agents in this model never completely abandon their hill-climbing specialization procedure. Some human subjects, foiled in early attempts, simply stop looking for welfare improvements altogether and settle into autarky. While agents may eventually become unwilling to trade, they can always discard this conservatism and reengage, and even when they employ a short-sighted learning rule, they never completely stop their search of the production space. Only when they are fully specialized and trading at the competitive equilibrium with no theft, do they actually settle into an unchangeable pattern of behavior.

[15] Kaufmann and Levin (1987), Kaufmann (1993), Altenberg (1997), Frenken (2006), Rivkin (2000), and Levinthal (1997) all explore hill-climbing processes on various "fitness landscapes" in the context of biology, organizational theory, and technological change. Figure 1 depicts a profit landscape faced by any *odd* agent trading with any *even* agent. The relationship between an agent's rate of specialization and its trading partner's rate of specialization is "epistatic" in that a change in one element bears heavily on the profitability (fitness) of the other under the consumption and trading algorithms that determine profit.

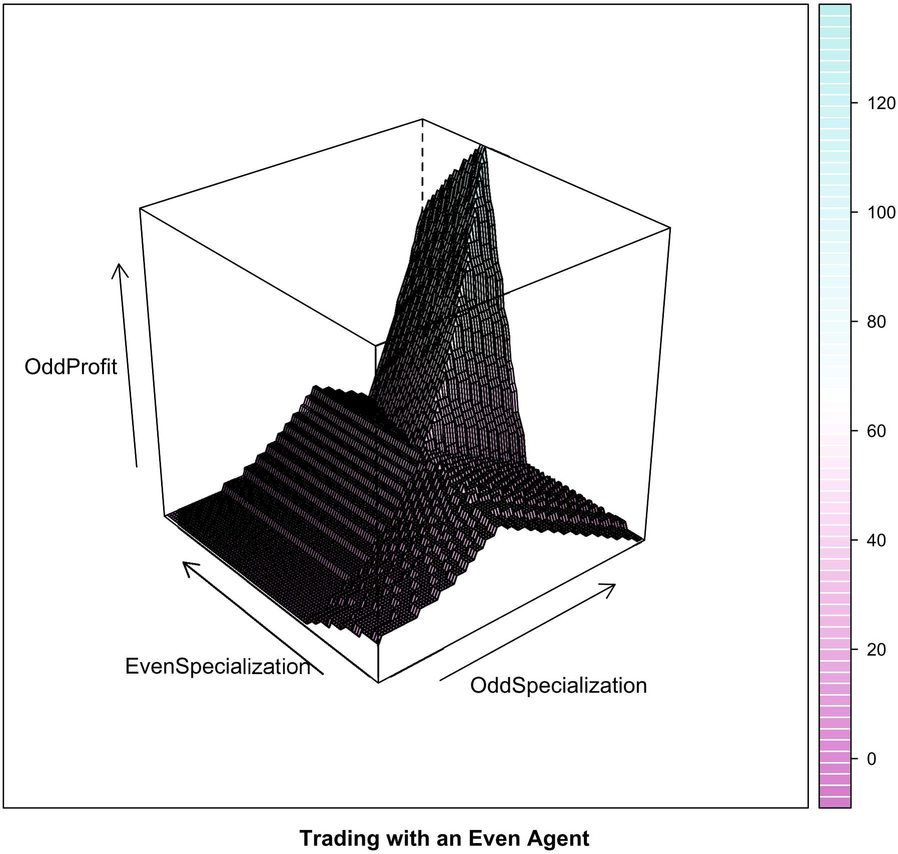

Fig. 1 Profit landscape for an *odd* agent

varying degrees depending on their conservatism, c, and probability of mimicry, m. Agents may decide to After loop 4 is completed, relevant data is recorded and the simulation begins a new period.

Parameter Update Pseudocode (continued)

Update Willingness to Trade of Each Agent i

If # of Trades > 0

$\quad$ *If* $w_i = unwilling$

$\quad\quad$ *If* $runif_{[0,1]} > c_i$

$\quad\quad$ Set $w_i = willing$

$\quad\quad$ *If* $\sum \Delta_i = 0$, *re-initialize* Δ_i

$\quad\quad$ *If* $runif_{[0,1]} > m_i$

$\quad\quad\quad$ Let $\psi = I \backslash i$

$\quad\quad\quad$ Choose non-autarkist $j \in \psi$, *s.t.* $e_{jt} > e_{it}$[16]

[16]This process is blind to the agent's type, so such decisions will not always be wise.

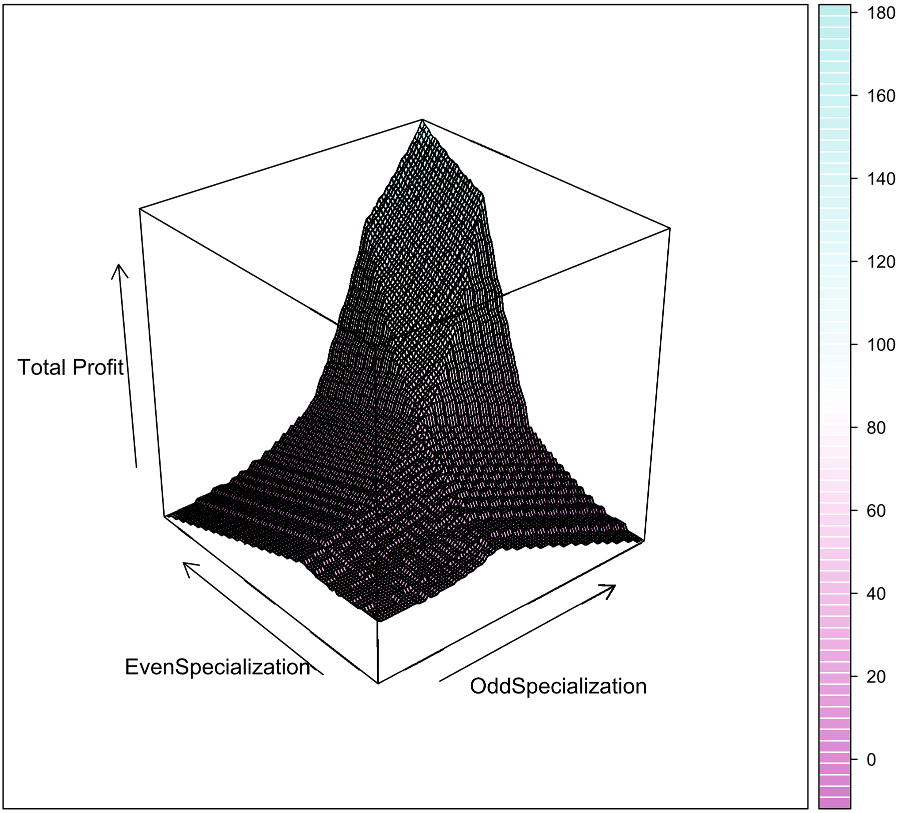

Fig. 2 Joint profit landscape: *odd* and *even*

$$\text{Set } \tau_i = \tau_j$$

End If $(\times 4)$

End Loop 4—Parameter Updates

A.7 Model parameterizations

I consider 6 different parameterizations, where each varies two values that guide agents' learning rules: (1) the level of earnings below which each agent chooses to adjust its learning rule when acting as an autarkist, or e^A, and (2) the level of earnings below which each agent chooses to adjust its learning rule when trading with another agent, or e^T. These two parameters are homogeneous in each session's agent population, exogenously imposed and, when combined with random draws that guide behavior, determine the evolution of the economy.

The treatment variables e^A and e^T act as proxies for an agent's risk aversion and/or discount rate; think of them as representing an agent's maximum acceptable risk at any given level of production and conditional on whether that agent acts in autarky or trades. A higher value of either means that agents will more quickly abandon incremental production-space searches that produce relative decreases in short-run earnings. Agents with lower values, on the other hand, will assume additional risk in the short run when exploring the behavior space in hopes of finding long-term benefits.

Simulation names refer directly to these values. Thus, the simulations with $e^T = 10$ and $e^A = 15$ will be referred to as the 1015 simulations. These values were chosen over time through a simple grid search process in order to acquire values that yield data nearly approximating that of the human data from CSW.

Appendix B: Experiment instructions

Page 1 This is an experiment in the economics of decision making. The instructions are simple, and if you follow them carefully and make good decisions you may earn a considerable amount of money which will be paid to you in CASH at the end of the experiment.

In this experiment you are **Person 2**. You and the other **7** people in this experiment each have the ability to produce two fictitious items: red and blue. For the first **10** seconds of each period, you will produce items in the upper left portion of your screen. Using the scroll bar, you can change the proportion of each second allocated to producing red and blue. Each person's production is displayed in the domino-shapes at the bottom of your screen.

When a domino-shape or house is selected, its contents are displayed in the top left portion of your screen. To select a domino-shape or house, left click on it and it will be become highlighted in *yellow*.

Page 2 After the production phase ends, the period continues for another **90** seconds. When the clock expires, you earn cash based upon the number of red and blue items that have been moved to your house. To select items to be moved, *left* click on an item or click on the red or blue buttons at the top of the screen. The yellow highlighted items can be moved by dragging with the *right* mouse button. The maximum number of *red* or *blue* items a house or field can hold is 170. (You cannot move items until the experiment has started or during the production phase.)

The specific information on how the *red* and *blue* items in your house generate earnings is given in the upper right corner of your screen. You personally earn (in cents) the minimum of the following two numbers:

2 *times number of red* items, number of blue items.

Or, think of it this way. You earn by consuming what's in your house in the proportion of 1 *red* to 2 *blue* items. *For every 1 unit of red you need 2 units of blue to earn 2 cents*. Your potential profit updates as items, unit by unit, are moved into your house.

Page 3 Everyone in this experiment can send text messages. Everyone can read all posted messages. In the center of the screen, you can type a message in the text box next to the send button. To send a text message press the **Send** button. There are two chat rooms. Messages sent to Chat Room A will only appear in chat room A. Message sent to Chat Room B will only appear in chat room B.

Under your house you can also post a one-line message that will be visible at all times to the other players.

You are free to discuss all aspects of the experiment, with the following exceptions: you may not reveal your name, discuss side payments, make threats, or engage in inappropriate language (including such shorthand as ‘WTF’). If you do, you will be excused and you will forfeit your earnings.

Page 4 During the experiment, every 7 periods will be a “break period” in which nobody produces anything but that the chat rooms are still open.

You can open a table of your production possibilities by clicking on the **Show %s** button. This table will fill in every time you change the proportion of time allocated using the scroll bar.

This is the end of the instructions. If you wish to explore how you produce red and blue items, click the **Practice** button. You may change the proportion of time allocated to producing red and blue items using the scroll bar, and you may **Practice** as many times as you wish. (You will not be able to move items until the experiment has begun.)

If you wish to review the instructions, you may go back at this time. If you feel you are prepared to proceed with the actual experiment, click on the **Start** button. The experiment will begin once everyone has clicked on the **Start** button. If you have a question that you feel was not adequately answered by the instructions, please raise your hand and ask the monitor before proceeding.

References

Altenberg, L. (1997). NK fitness landscapes. In T. Back, D. Fogel, & Z. Michalewicz (Eds.), *The handbook of evolutionary computation*. Oxford: Oxford University Press.

Frenken, K. (2006). A fitness landscape approach to technological complexity, modularity, and vertical disintegration. *Structural Change and Economic Dynamics*, *17*, 288–305.

Gale, D., & Shapley, L. S. (1962). College admissions and the stability of marriage. *American Mathematical Monthly*, *69*(1), 9–14.

Kaufmann, S.A. (1993). *The origins of order: self-organization and selection in evolution*. New York: Oxford University Press.

Kaufmann, S., & Levin, S. (1987). Toward a general theory of adaptive walks on rugged landscapes. *Journal of Theoretical Biology*, *128*, 11–45.

Kimbrough, E. O. (forthcoming). Heuristic learning and the discovery of specialization and exchange. *Journal of Economic Dynamics and Exchange*.

Levinthal, D. A. (1997). Adaptation on rugged landscapes. *Management Science*, *43*(7), 934–950.

Rivkin, J. W. (2000). Imitation of complex strategies. *Management Science*, *46*(6), 824–844.

Skyrms, B. (2004). *The stag hunt and the evolution of social structure*. Cambridge: Cambridge University Press.

Starr, R. M. (1976). Decentralized nonmonetary trade. *Econometrica*, *44*(5), 1087–1089.

Tesfatsion (1997). A trade network game with endogenous partner selection. In H. Amman, B. Rustem, & A. B. Whinston (Eds.), *Computational approaches to economic problems: advances in computational economics*. New York: Springer.