Online Appendix: Coarse Payoff-Assessment Learning Model

1 Algorithm for transforming $\mathcal{T}$ into $\mathcal{T}'$

We extend the state-dependent payoff function to the set of similarity classes. Using the same notation as before, we define $\mu_\psi : \mathcal{S}_\psi \to \mathbb{R}$, meaning that Alice derives an expected payoff of $\mu_\psi(s)$ for picking an alternative from the similarity class $s \in \mathcal{S}_\psi$ in state ψ. If a similarity class s contains a single alternative $c \in s$ available at state ψ, then $\mu_\psi(s) = \mu_\psi(c)$. However, if multiple alternatives ($n > 1$) from within a similarity class s are available in state ψ, the expected payoff for the class s in that state is simply the average of the payoffs of all alternatives in s available at ψ, i.e., $\mu_\psi(s) = \frac{1}{n}\sum_{c \in s} \mu_\psi(c)$. This is because we assume that Alice cannot distinguish between alternatives within the same similarity class, leading her to randomize uniformly among them when multiple such alternatives are available at a given node.[1] We notice that applying the similarity transformation may result in several nodes presenting indistinguishable choice problems over identical subsets of similarity classes from Alice's perspective. This can include nodes offering trivial unary choice problems, where Alice can only choose among alternatives within a single similarity class. Naturally, Alice would collapse all such nodes into a single representative node. In light of this observation, we construct a new decision tree $\mathcal{T}'$, that modifies the original decision tree $\mathcal{T}$ without altering Alice's underlying choice incentives but crucially eliminates redundant nodes. In $\mathcal{T}'$, as usual, nature moves at the root $\mathbf{r}'$ and randomly draws a node $\omega \in \mathbf{\Omega}$ according to a fixed probability mass function $\mathbf{p}(\omega)$. We provide an algorithm for constructing $\mathcal{T}'(\mathbf{\Omega}, \mathbf{p}, \pi)$ from $\mathcal{T}$.

Condition 1 (No choice-redundancy)**.** *For a node* $\psi \in \mathbf{\Psi}$*,* $\nexists\ \chi \in \mathbf{\Psi}$ *s.t.* $\psi \neq \chi \wedge \mathcal{S}_\psi = \mathcal{S}_\chi$.

We initialize $\mathcal{T}'$ with an empty set of nodes denoted by $\mathbf{\Omega}$. We draw a node ν in $\mathbf{\Psi}$ from $\mathcal{T}$ without replacement. ν either satisfies or violates Condition 1.

Case I $\nu \in \mathbf{\Psi}$ satisfies Condition 1. Then, $\nu \in \mathbf{\Omega}$. Also, $\mathbf{p}(\nu) = \mathbf{f}(\nu)$, $\mathcal{S}_\nu^{\mathcal{T}'} = \mathcal{S}_\nu^{\mathcal{T}}$, and $\pi_\nu = \mu_\nu$. Any node in $\mathbf{\Psi}$ that satisfies Condition 1 is immediately added to $\mathbf{\Omega}$. Such a node in $\mathcal{T}$ is one that presents Alice with a distinct choice problem. It's essentially duplicated onto $\mathcal{T}'$ with the same probability of being reached by nature as in $\mathcal{T}$. The set of options (similarity classes) available to Alice at such a node in $\mathcal{T}'$ is the same as in $\mathcal{T}$. Also, the payoff that Alice receives by choosing an alternative in a similarity class that is available at node ν in $\mathcal{T}'$ is equal to the payoff she receives by choosing the same alternative at node ν in $\mathcal{T}$. Finally, we remove ν from $\mathbf{\Psi}$, i.e., $\nu \notin \mathbf{\Psi}$.

[1]For e.g., consider an agent, Ted, who must choose between driving a car and taking a blue bus to work. Initially indifferent between the two options, he assigns equal valuations to both and chooses each with a 50% probability. Now, if a red bus is introduced on the same route, Ted perceives the blue and red buses (that differ only in color) as similar based on average commute times, grouping them into a single similarity class — bus. In this case, Ted remains equally likely to choose either the car or the bus. Within the bus category, he is indifferent between the blue and red buses, choosing each with equal probability whenever he opts for the bus. This example highlights that our model is more aligned with nested logit models in econometrics, rather than multinomial logit models, by accommodating existing similarities among alternatives.

Case II $\nu \in \mathbf{\Psi}$ violates Condition 1. That is, $\exists\ \chi \in \mathbf{\Psi}$ s.t. $\nu \neq \chi \wedge \mathcal{S}_\nu = \mathcal{S}_\chi$. Given such a node ν, we find all possible nodes $\chi \in \mathbf{\Psi}$ that are likewise choice-equivalent to ν according to Condition 1 and add them to the set $\mathcal{I}$ along with ν. We know that $|\mathcal{I}|$ is an integer greater than 1. We consolidate the nodes in $\mathcal{I}$ into one terminal node $\omega \in \mathbf{\Omega}$ in the decision tree $\mathcal{T}'$ such that $\mathcal{S}_\omega^{\mathcal{T}'} = \mathcal{S}_\nu^{\mathcal{T}}$ and $\mathbf{p}(\omega) = \sum_{i \in \mathcal{I}} \mathbf{f}(i)$.

$$\pi_\omega(s) = \frac{\sum_{i \in \mathcal{I}} \mathbf{f}(i) \mu_i(s)}{\sum_{i \in \mathcal{I}} \mathbf{f}(i)},$$

represents the expected payoff that Alice receives by choosing an alternative in a similarity class $s \in \mathcal{S}_\omega$ at node ω in tree $\mathcal{T}'$. Finally, we remove the set $\mathcal{I}$ from $\mathbf{\Psi}$, i.e., $\mathcal{I} \not\subseteq \mathbf{\Psi}$.

The procedure described above is repeated for every remaining node in $\mathbf{\Psi}$ until $\mathbf{\Psi} = \emptyset$. Thus, nodes in $\mathbf{\Psi}$ satisfying Condition 1 are mapped one-to-one from $\mathcal{T}$ to $\mathcal{T}'$ while nodes violating Condition 1 in $\mathcal{T}$ are aggregated into a representative node that is then added to $\mathcal{T}'$. It's easy to see that for a given node $\omega \in \mathbf{\Omega}$ in $\mathcal{T}'$, $\nexists\ \chi \in \mathbf{\Omega}$ s.t. $\omega \neq \chi \wedge \mathcal{S}_\omega = \mathcal{S}_\chi$. In the remainder of this paper, we will work directly with the similarity-transformed decision tree $\mathcal{T}'$. Finally, $\mathcal{S} = \cup_{\omega \in \mathbf{\Omega}} \mathcal{S}_\omega$ denotes the grand set of similarity classes available to Alice in the decision tree $\mathcal{T}'$.

2 Proof of Proposition 1

Proof. We've already established the existence and the global asymptotic stability of the unique smooth valuation equilibrium that arises in the neighborhood of a mixed valuation for a sufficiently large sensitivity parameter in the proof of Theorem 4. We're left with having to show that the mixed valuation equilibrium is in fact a fully-mixed VE and not a partially-mixed VE.

Recall that, in any partially-mixed VE, at least one of the following two conditions must hold: there exists a class k such that either $v_k^* < v_i^* = v_j^*$ or $v_k^* > v_i^* = v_j^*$, in equilibrium. Now, following the same reasoning as in the proof of Theorem 4, a sufficiently large constant z added to the trivial choice payoffs can also rule out the set of partially-mixed valuation equilibria in a decision tree that satisfies Assumption 5.1. The underlying logic remains the same: the higher the equilibrium valuation of a similarity class, the more choice nodes it is selected at, resulting in a smaller weight on z in the consistency calculation. By making z sufficiently large, one can reverse the order relation among the equilibrium valuations leading to a contradiction.

To begin, let's assume $v_k^* < v_i^* = v_j^*$. In the consistency calculation, the weight on z in v_k^* would be larger than the weight on z in v_i^* and v_j^*. This is because, at the very least, k would be chosen with probability 0 at the binary choice nodes $\{i, k\}$ and $\{j, k\}$, as well as at the tertiary choice node $\{i, j, k\}$. By making z sufficiently large, one can ensure $v_k^* \geq v_i^* = v_j^*$, arriving at a contradiction. Conversely, let's assume $v_k^* > v_i^* = v_j^*$. In the consistency calculation, the weight on z in v_k^* would be smaller than the weight on z in v_i^* and v_j^*. This is because, at the very least, k would be chosen with probability 1 at the binary choice nodes $\{i, k\}$ and $\{j, k\}$, as well as at the tertiary choice node $\{i, j, k\}$. Again, by making z sufficiently large, one can ensure $v_k^* \leq v_i^* = v_j^*$, leading to a contradiction. Thus, we've ruled out both strict pure VE and partially-mixed VE implying the existence of a fully-mixed VE.

Therefore, invoking Lemma 2.1, there exist a $\hat{z} > 0$ and a $\hat{\beta} > 0$ such that $\forall\ z > \hat{z}$ and $\forall\ \beta > \hat{\beta}$, the unique SVE lies in the neighborhood of a fully-mixed VE where the valuations of all the similarity classes in $\mathcal{S}$ are identical in the high-sensitivity limit.[2] □

[2] We note that this result stands in stark contrast with the fact that when an agent has $n > 2$ similarity classes available to her, the set of fully-mixed valuation equilibria in a decision tree with generic payoffs lies on a manifold. For e.g., with just 3 similarity classes, the set of fully-mixed VE is characterized by a polynomial system with 2 equations (indifference conditions) and 5 unknowns (mixing probabilities at the non-trivial choice nodes). The existence of VE result in **?** establishes that there exists at least 1 mixed VE (once we've ruled out the strict pure VE and the partially-mixed VE) implying that there are infinite solutions. Thus, the set of mixed VE lies on a manifold of dimension at least 3 while the CPAL dynamics converge to a unique mixed SVE even in the high-sensitivity limit.

3 Additional Illustrations

3.1 Example: The Good, the Bad, and the Unsteady

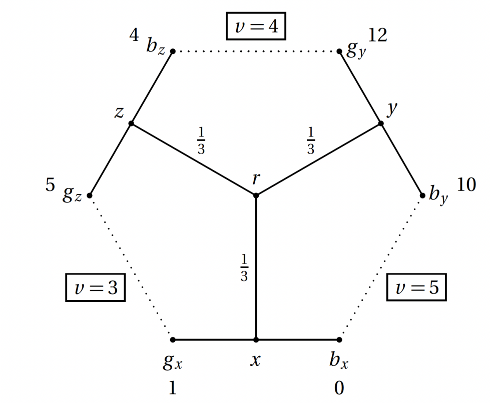

Figure 1: Good/Bad Decisions

The decision tree is depicted by the solid lines in Figure 1. At the root r, nature chooses one of three nodes x, y, and z with equal probability. At each of these nodes, the decision maker can choose between a good alternative and a bad alternative, where the former has a higher payoff. The payoffs are written next to these alternatives. The three dotted lines connect similar nodes. Thus, the set of alternatives is partitioned into the similarity classes $i = \{g_x, g_z\}$, $j = \{b_x, b_y\}$ and $k = \{g_y, b_z\}$. We find this example interesting because it illustrates a counter-intuitive case where as a result of the similarity grouping, making the worst decision at every node is a valuation equilibrium (VE) while making the best one is not. The unique valuation that is consistent with the strategy that selects the bad alternative at every node is $(v_i, v_j, v_k) = (3, 5, 4)$.

It turns out, however, that the worst-decision VE does not correspond to a rest-point of the CPAL model in this decision tree even in the high-sensitivity limit. To show this, we derive the RHS of the CPAL dynamics (Eq. 6) given the primitives of this decision tree, evaluate it at the valuation $(3, 5, 4)$ for an arbitrary, finite β and then compute the limit of the expression as $\beta \uparrow \infty$. For the sake of brevity, we choose not to include the calculations here but it can be verified that the solution to our exercise above is $(2, 0, 0)$. At steady-state, we'd expect the RHS to be $(0, 0, 0)$. Evidently, the worst-decision VE is not a steady-state of the CPAL model. Specifically, at the worst-decision VE, it is the valuation of the similarity class $i = \{g_x, g_z\}$ that has not yet reached stationarity.

We conclude that while the worst-decision VE manages to meet the optimality condition of valuation equilibrium, it fails to do so in the smoother CPAL model even in the exact best-response limit. Therefore, while a sequence of β-dependent steady-states of the CPAL model converges to a VE as $\beta \uparrow \infty$, the converse is not true. That is, there exist VE that cannot be characterized as the limit points of the steady-states of CPAL dynamics. In fact, the set of Smooth Valuation Equilibria in the high sensitivity limit characterizes a refinement (subset) of the set of Valuation Equilibria.

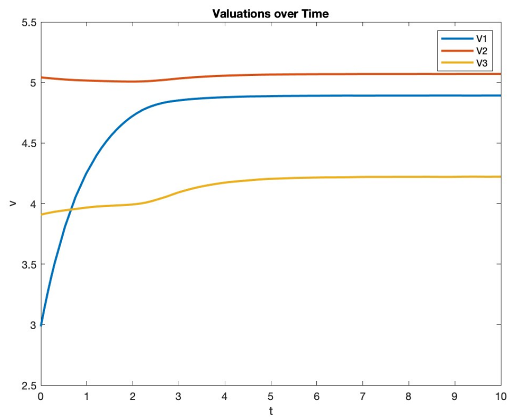

Figure 2: CPAL dynamics with initial valuation $(3, 5, 4)$ & $\beta = 50$

We also run numerical simulations to confirm our claim. Indeed, starting with an initial valuation of $(3, 5, 4)$ and using a large sensitivity parameter, we observe in Figure 2 that the valuation system rapidly moves on to a different point before coming to rest. In fact, it is v_i that does not stay at rest at $v_i = 3$ while v_j and v_k seem to be at rest at 5 and 4 respectively, thereby reinforcing our insight from the theoretical exercise above. We repeat the simulations a large number of times each time varying the initial conditions while using a large enough β. We observe convergence to two distinct rest-points in the long-run: one at $(v_i \approx 5, v_j = 5, v_k \approx 4)$ that corresponds to a partially-mixed VE as seen in Figure 2 and the other at $(v_i = 1, v_j = 0, v_k = 8)$ that corresponds to a strict pure VE as seen in Figure 3.

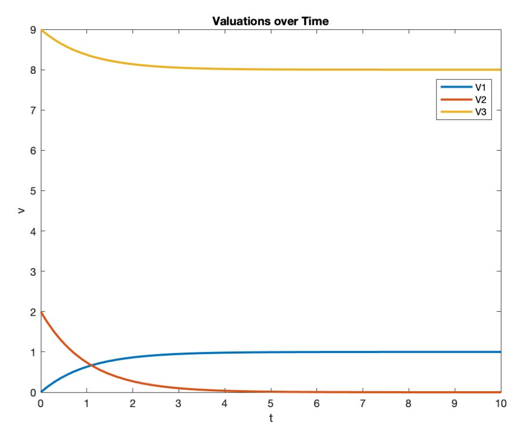

Figure 3: CPAL dynamics with initial valuation $(0, 2, 9)$ & $\beta = 50$

We linearize Equation 6 around a given steady-state valuation and compute the eigenvalues of the corresponding Jacobian matrix for an arbitrary β and evaluate their limits as $\beta \uparrow \infty$. We conduct this exercise for both the steady-state valuation systems characterized above. We find that the eigenvalues of the Jacobian matrix are negative for both the rest-points, as $\beta \uparrow \infty$, confirming their local asymptotic stability as observed in the numerical simulations.

3.2 Example: The Road Not Taken

The decision tree is depicted by the solid lines in Fig. 4. At the root r, nature chooses one of two nodes x and y with equal probability. At node x, the decision-maker can choose between alternatives m_x and l_x, and at node y, between m_y and r_y. The dotted line connects similar alternatives. The set of alternatives is partitioned into three similarity classes, $L = \{l_x\}$, $M = \{m_x, m_y\}$ and $R = \{r_y\}$. There are a total of three valuation equilibria - two pure and one partially-mixed.

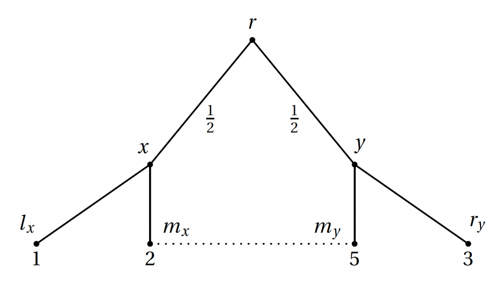

Figure 4: The Road Not Taken

- The strategy that selects the alternative in M at each of the nodes x and y is a strict pure VE. The corresponding valuation vector is $(v_L = 1, v_M = 3.5, v_R = 3)$ and the strategy is clearly optimal for this valuation. We verify, by direct computation, that the valuation $(1, 3.5, 3)$ is a steady-state of the CPAL model as $\beta \uparrow \infty$. The numerical simulation seen in Fig. 5 points to the same. Also, with a large enough sensitivity parameter, we see strong evidence of convergence to the steady-state starting from a nearby initial valuation system.

- The strategy that selects the alternatives in M at node x and R at y is also a strict pure VE. The corresponding valuation is $(v_L = 1, v_M = 2, v_R = 3)$ and the strategy is clearly optimal for this valuation. We verify, by direct computation, that the valuation $(1, 2, 3)$ is a steady-state of the CPAL model as $\beta \uparrow \infty$. The numerical simulation seen in Fig. 6 points to the same. Also, with a large sensitivity parameter, we see evidence of convergence to the steady-state starting from a nearby initial valuation system.

- The strategy that always selects the alternative in M at node x and uniformly randomizes between the alternatives in M and R at y is a partially-mixed VE. The corresponding valuation is $(v_L = 1, v_M = 3, v_R = 3)$ and the strategy is clearly optimal for this valuation. We verify, by direct computation, that the valuation $(1, 3, 3)$ is a steady-state of the CPAL model as $\beta \uparrow \infty$. However, from the numerical simulation presented in Fig. 7, it seems that the steady-state is unstable. A small perturbation to an initial valuation around $(1, 3, 3)$ causes the valuation system to move away and eventually come to rest at one of the strict pure VE. Indeed, when we linearize the system around $(1, 3, 3)$ and compute the eigenvalues of the corresponding Jacobian matrix in the $\beta \uparrow \infty$ limit, we find that one of the eigenvalues is positive. Thus, this

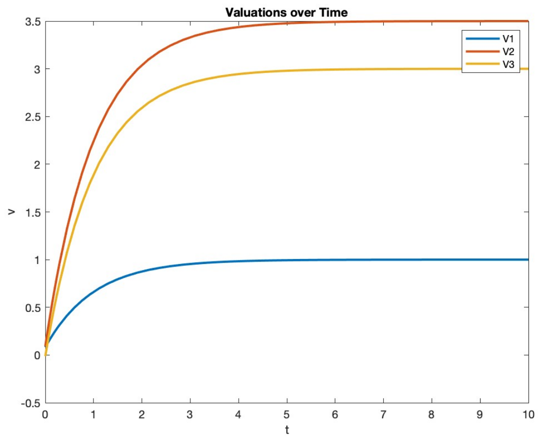

Figure 5: strict pure VE at $(1, 3.5, 3)$; $\beta = 50$

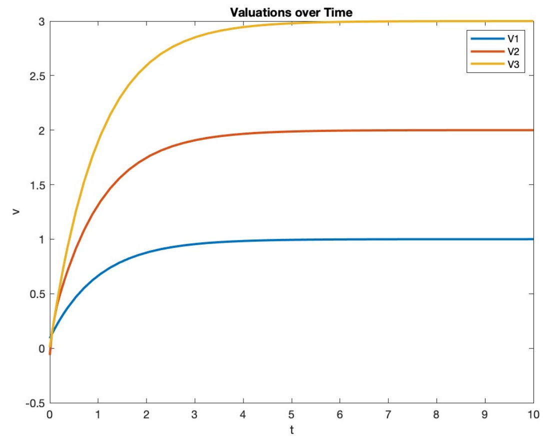

Figure 6: strict pure VE at $(1, 2, 3)$; $\beta = 50$

partially-mixed VE is an asymptotically unstable rest-point of the CPAL model.

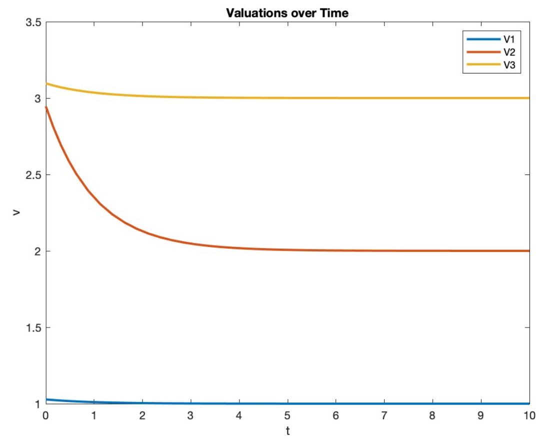

Figure 7: Asymptotically Unstable Partially-Mixed VE at $(1, 3, 3)$; $\beta = 50$

This example is fascinating for two key reasons. First, it demonstrates the possibility of multiple equilibria in a decision problem. Second, it allows us to characterize at least one asymptotically unstable equilibrium of the CPAL model.

4 Supplementary Material

4.1 Linearization Theorem

Consider an autonomous non-linear dynamical system:

$$\dot{x} = f(x)$$

with $f : \mathbb{R}^n \to \mathbb{R}^n$ being continuously differentiable, and suppose x^* is an equilibrium point, i.e., $f(x^*) = 0$. The system can be linearized around the equilibrium point x^* as:

$$\dot{\delta x} = A \delta x$$

where $\delta x = x - x^*$ and A is the Jacobian matrix of f evaluated at x^*:

$$A = \left. \frac{\partial f}{\partial x} \right|_{x = x^*}$$

The Linearization Theorem (or (Hartman-Grobman Theorem) states:

1. If all eigenvalues of A have negative real parts, then the equilibrium x^* of the non-linear system is locally asymptotically stable.

2. If any eigenvalue of A has a positive real part, then the equilibrium x^* of the non-linear system is unstable.

3. If A has eigenvalues with zero real parts and no eigenvalues with positive real parts, the theorem does not provide any conclusion about the stability of the non-hyperbolic equilibrium x^* in the non-linear system. In these cases, higher-order terms or other methods such as Lyapunov functions might be needed to determine stability.

While the Linearization Theorem is powerful, it has its limitations in the sense that it only provides information about local behavior around a hyperbolic equilibrium. It says nothing about the global behavior of the system.

Asymptotic Stability

Let $\mathcal{D}$ be an open set in $\mathbb{R}^n$ containing the origin, and let $x^* = 0$ be an equilibrium point for the dynamical system

$$\dot{x} = f(x)$$

where $f : \mathcal{D} \to \mathbb{R}^n$ is a continuous vector field. The equilibrium $x^* = 0$ is said to be **locally asymptotically stable** if it satisfies both of the following conditions:

1. **Lyapunov Stability**: For every $\varepsilon > 0$, there exists a $\delta(\varepsilon) > 0$ such that, if the initial condition $x(0)$ satisfies $\|x(0)\| < \delta$, then for all $t \geq 0$, $\|x(t)\| < \varepsilon$.

2. **Convergence**: There exists $\delta > 0$ such that, if the initial condition $x(0)$ satisfies $\|x(0)\| < \delta$, then $\lim_{t \to \infty} x(t) = 0$.

Additionally, if this property holds for all initial conditions $x(0) \in \mathbb{R}^n$, then the equilibrium x^* is said to be **globally asymptotically stable**.

4.2 Gershgorin Circle Theorem

Let A be an $n \times n$ matrix with complex entries a_{ij}. For each $i \in \{1, 2, \ldots, n\}$, define the ith Gershgorin circle C_i in the complex plane as:

$$C_i = \left\{ z \in \mathbb{C} : |z - a_{ii}| \leq \sum_{\substack{j=1 \\ j \neq i}}^{n} |a_{ij}| \right\}$$

Then, every eigenvalue of A lies within at least one of the Gershgorin circles C_i.

Proof. Let β be an eigenvalue of A and $\mathbf{x}$ be a corresponding eigenvector, where $\mathbf{x} \neq \mathbf{0}$. Without loss of generality, assume that among the entries of $\mathbf{x}$, x_k has the largest magnitude, i.e., $|x_k| = \max\{|x_1|, |x_2|, \ldots, |x_n|\}$, and $|x_k| > 0$ since $\mathbf{x}$ is not the zero vector.

Since $A\mathbf{x} = \beta\mathbf{x}$, we have for the kth entry:

$$\sum_{j=1}^{n} a_{kj} x_j = \beta x_k$$

Rearranging terms gives:

$$a_{kk} x_k + \sum_{\substack{j=1 \\ j \neq k}}^{n} a_{kj} x_j = \beta x_k$$

Subtracting $a_{kk} x_k$ from both sides and then dividing by x_k (which is nonzero) yields:

$$\beta - a_{kk} = \sum_{\substack{j=1 \\ j \neq k}}^{n} a_{kj} \frac{x_j}{x_k}$$

Taking the absolute value of both sides, we obtain:

$$|\beta - a_{kk}| = \left| \sum_{\substack{j=1 \\ j \neq k}}^{n} a_{kj} \frac{x_j}{x_k} \right| \leq \sum_{\substack{j=1 \\ j \neq k}}^{n} |a_{kj}| \left| \frac{x_j}{x_k} \right|$$

Since $|x_k| \geq |x_j|$ for all j, we have $\left| \frac{x_j}{x_k} \right| \leq 1$, and therefore:

$$|\beta - a_{kk}| \leq \sum_{\substack{j=1 \\ j \neq k}}^{n} |a_{kj}|$$

β lies within the circle C_k, as $|z - a_{kk}| \leq \sum_{j \neq k} |a_{kj}|$ defines the circle in the complex plane where $z = \beta$. Thus, every eigenvalue of A must lie within at least one of the Gershgorin circles C_i. $\square$