Robust Expectations Adaptation

Te Bao[1] Peter Bossaerts[2] **Jiaoying Pei**[2]

[1]Nanyang Technological University, Singapore [2]University of Cambridge, U.K.

May 2026

Problem Statement

- We want to understand how predictions are made by real humans.

Problem Statement

- We want to understand how predictions are made by real humans.
- This has wide varieties of implications. In economics, one include understanding price dynamics; bubbles/crashes, as stock price is self-referential.

What has the literature proposed and why it fails: part 1

- The traditional model for learning to forecast in economics is *adaptive expectations*, aka as *Rescorla Wagner Rule* or *Prediction-Error Learning* in psychology, neuroscience and computer science.

$$p^e_{t+1} = \underbrace{p^e_t}_{\text{forecast of observed price } p_t,\text{ indexed } t=1,2,\ldots} + \overbrace{\alpha}^{\text{learning rate}} \underbrace{(p_t - p^e_t)}_{\text{prediction error } = \epsilon_t},$$

What has the literature proposed and why it fails: part 1

- The traditional model for learning to forecast in economics is *adaptive expectations*, aka as *Rescorla Wagner Rule* or *Prediction-Error Learning* in psychology, neuroscience and computer science.

$$p^e_{t+1} = \underbrace{p^e_t}_{\text{forecast of observed price } p_t, \text{ indexed } t=1,2,\ldots} + \overbrace{\alpha}^{\text{learning rate}} \underbrace{(p_t - p^e_t)}_{\text{prediction error } = \epsilon_t},$$

- The learning rate α (of how much you learn from your mistakes) is assumed to be constant, e.g., Kalman filter in Bordalo et al. (2020).

What has the literature proposed and why it fails: part 1

- The traditional model for learning to forecast in economics is *adaptive expectations*, aka as *Rescorla Wagner Rule* or *Prediction-Error Learning* in psychology, neuroscience and computer science.

$$p^e_{t+1} = \underbrace{p^e_t}_{\text{forecast of observed price } p_t,\text{ indexed } t=1,2,\ldots} + \overbrace{\alpha}^{\text{learning rate}} \underbrace{(p_t - p^e_t)}_{\text{prediction error } = \epsilon_t},$$

- The learning rate α (of how much you learn from your mistakes) is assumed to be constant, e.g., Kalman filter in Bordalo et al. (2020).
- But a constant α cannot be optimal in economics b/c the object to be forecasted is endogenous. That is, the system forecast-observation is *self-referential*: the target moves, among others, driven by the forecasts: Marcet and Sargent.

What has the literature proposed and why it fails: part 1

- The traditional model for learning to forecast in economics is *adaptive expectations*, aka as *Rescorla Wagner Rule* or *Prediction-Error Learning* in psychology, neuroscience and computer science.

$$p^e_{t+1} = \underbrace{p^e_t}_{\text{forecast of observed price } p_t, \text{ indexed } t=1,2,\ldots} + \overbrace{\alpha}^{\text{learning rate}} \underbrace{(p_t - p^e_t)}_{\text{prediction error } = \epsilon_t},$$

- The learning rate α (of how much you learn from your mistakes) is assumed to be constant, e.g., Kalman filter in Bordalo et al. (2020).
- But a constant α cannot be optimal in economics b/c the object to be forecasted is endogenous. That is, the system forecast-observation is *self-referential*: the target moves, among others, driven by the forecasts: Marcet and Sargent.
- Self-referentiality leads to *non-stationarity*: *the parameters of the data generating process (DGP) of the objects* to be forecasted changes. Non-stationarity is one of the reasons why forecasting models in, e.g., financial markets, have poor out-of-sample performance. Bossaerts and Hillion

- Nonstationarities: when mean and variance of price changes over times.

- Nonstationarities: when mean and variance of price changes over times.
- Mean changes: regime shift. Oil price. α cannot be a constant.

- Nonstationarities: when mean and variance of price changes over times.
- Mean changes: regime shift. Oil price. α cannot be a constant.
- Solution? In psychology and neuroscience, robustness is introduced by allowing α to increase with *surprise*, i.e., with $|\epsilon_t|$: *Pearce-Hall Prediction-Error Learning*. Pearce and Hall.

- Nonstationarities: when mean and variance of price changes over times.
- Mean changes: regime shift. Oil price. α cannot be a constant.
- Solution? In psychology and neuroscience, robustness is introduced by allowing α to increase with *surprise*, i.e., with $|\epsilon_t|$: *Pearce-Hall Prediction-Error Learning*. Pearce and Hall.
- Maladaptive if outliers revert as in d'Acremont and Bossaerts (&2/3 outliers revert in real financial market). Also used in economics: Carvalho et al. (2023). Example: trade war

What has the literature proposed and why it fails: part 2

- In general, non-stationarity may affect variability of object to be forecasted (i.e., some periods to have big swings and others to be calm).

What has the literature proposed and why it fails: part 2

- In general, non-stationarity may affect variability of object to be forecasted (i.e., some periods to have big swings and others to be calm).
- One way is to track whether there are "regime switches" (Hamilton; switchpoint detection in engineering) and apply whichever rule (in the form of a new learning rate α) is best; *Heuristic Switching Model*: Anufriev and Hommes. The idea is to focus on switchpoint detection. This requires the forecaster to know beforehand what the possible "regimes" could be, and that there are sufficiently long, stable (stationary) episodes. If not, problems with convergence: Csáji and Monostori. Example: electricity price predictions during COVID.
- In economics, the standard way to *robust*ly deal with nonstationarity is to prepare for the worst. Hansen and Sargent's Min Max. This is obviously not ideal, if only because the worst possible is *unknown* & *rarely happen*.

Proposed solution

- Here, we study how humans should and do adapt expectations to account for non-stationarity.

Proposed solution

- Here, we study how humans should and do adapt expectations to account for non-stationarity.
- We propose that you need to learn from your mistakes (quickly and efficiently) when regime shift (i.e., outliers do not revert). But keep the existing learning rate when it's within expectation.

Proposed solution

- Here, we study how humans should and do adapt expectations to account for non-stationarity.
- We propose that you need to learn from your mistakes (quickly and efficiently) when regime shift (i.e., outliers do not revert). But keep the existing learning rate when it's within expectation.
- A supervisory system determines whether this is a regime shift or not **on the fly**, using a reference model of *(fixed-parameter) Kalman Filter Model* (Bordalo et al., 2020).

Proposed solution

- Here, we study how humans should and do adapt expectations to account for non-stationarity.
- We propose that you need to learn from your mistakes (quickly and efficiently) when regime shift (i.e., outliers do not revert). But keep the existing learning rate when it's within expectation.
- A supervisory system determines whether this is a regime shift or not **on the fly**, using a reference model of *(fixed-parameter) Kalman Filter Model* (Bordalo et al., 2020).
- The approach is a form of satisficing (Simon, 1955), since the agent is not trying to estimate the best model for the situation at hand, because such estimation would take too much time (the model changes) or the parameters (that would capture the nonstationarity) is prohibitively high.

Why MRAC?

- Supervisory role in decision making.

Why MRAC?

- Supervisory role in decision making.
- Anterior insula is found to activate across different cognitive tasks.

Why MRAC?

- Supervisory role in decision making.
- Anterior insula is found to activate across different cognitive tasks.
- While other parts of the brain mostly activate for certain task, anterior insula is well connected, receive inputs from lower level brain (sensory) and also communicate with all the other parts of the brain.

Why MRAC?

- Supervisory role in decision making.
- Anterior insula is found to activate across different cognitive tasks.
- While other parts of the brain mostly activate for certain task, anterior insula is well connected, receive inputs from lower level brain (sensory) and also communicate with all the other parts of the brain.
- Small part of the brain and intuitively, cannot be doing everything.

Why MRAC?

- Supervisory role in decision making.
- Anterior insula is found to activate across different cognitive tasks.
- While other parts of the brain mostly activate for certain task, anterior insula is well connected, receive inputs from lower level brain (sensory) and also communicate with all the other parts of the brain.
- Small part of the brain and intuitively, cannot be doing everything.
- Reactive for reward prediction error (surprise, or size of the prediction error). Not reacting to reward. Seems to be a good position to supervise performance.

Findings: Theory

- With the Kalman filter as a reference model, the prescription from Robust Expectations Adaptation is to *increase the learning rate as a function fo prediction error auto-covariance*. Using only 2 periods to estimate this autocovariance, this means:

$$\alpha_t \propto \varepsilon_t \varepsilon_{t-1}$$

- Intuition? Positive auto-covariance suggests under-reaction. Negative auto-covariance suggest over-reaction.
- The rule is the analogous to the Delta-Bar-Delta Algorithm in Machine Learning (Sutton, 1992).
- Robust control can interfere with (antagnize) attempts to apply optimal control (as above), to prevent overchasing the noise. When it does, it is uncovered that anterior insula is engaged (d'Acremont and Bossaerts, 2016).

Findings: Empirical

- We study 40,000+ *price forecasts* in a series of experiments on self-referentiality in markets. The experiments allow us to **enable self-referentiality and its nature**. (Bao et al., 2012, 2013, 2017; Bao and Hommes, 2019; Bao et al., 2024)

Findings: Empirical

- We study 40,000+ *price forecasts* in a series of experiments on self-referentiality in markets. The experiments allow us to **enable self-referentiality and its nature**. (Bao et al., 2012, 2013, 2017; Bao and Hommes, 2019; Bao et al., 2024)
- When imposing a constant learning rate α, we replicate the result in Bordalo ea that it is too high in positive feedback markets (e.g., a stock market). Even so, α needs to be higher than observed to reach optimality.

Findings: Empirical

- We study 40,000+ *price forecasts* in a series of experiments on self-referentiality in markets. The experiments allow us to **enable self-referentiality and its nature**. (Bao et al., 2012, 2013, 2017; Bao and Hommes, 2019; Bao et al., 2024)
- When imposing a constant learning rate α, we replicate the result in Bordalo ea that it is too high in positive feedback markets (e.g., a stock market). Even so, α needs to be higher than observed to reach optimality.
- We study estimates of changes in the learning rate: their *direction* (up/down).

Findings: Empirical

- We study 40,000+ *price forecasts* in a series of experiments on self-referentiality in markets. The experiments allow us to **enable self-referentiality and its nature**. (Bao et al., 2012, 2013, 2017; Bao and Hommes, 2019; Bao et al., 2024)
- When imposing a constant learning rate α, we replicate the result in Bordalo ea that it is too high in positive feedback markets (e.g., a stock market). Even so, α needs to be higher than observed to reach optimality.
- We study estimates of changes in the learning rate: their *direction* (up/down).
- We reject the Pearce-Hall Prediction-Error Learning: learning rate changes are insensitive to the surprise ($|\epsilon_t|$).

Findings: Empirical

- We study 40,000+ *price forecasts* in a series of experiments on self-referentiality in markets. The experiments allow us to **enable self-referentiality and its nature**. (Bao et al., 2012, 2013, 2017; Bao and Hommes, 2019; Bao et al., 2024)
- When imposing a constant learning rate α, we replicate the result in Bordalo ea that it is too high in positive feedback markets (e.g., a stock market). Even so, α needs to be higher than observed to reach optimality.
- We study estimates of changes in the learning rate: their *direction* (up/down).
- We reject the Pearce-Hall Prediction-Error Learning: learning rate changes are insensitive to the surprise ($|\epsilon_t|$).
- We find convincing evidence that the learning rate is modulated by the sign of the prediction error autocovariance ($\epsilon_t \epsilon_{t-1}$)...

Findings: Empirical

- We study 40,000$+$ *price forecasts* in a series of experiments on self-referentiality in markets. The experiments allow us to **enable self-referentiality and its nature**. (Bao et al., 2012, 2013, 2017; Bao and Hommes, 2019; Bao et al., 2024)
- When imposing a constant learning rate α, we replicate the result in Bordalo ea that it is too high in positive feedback markets (e.g., a stock market). Even so, α needs to be higher than observed to reach optimality.
- We study estimates of changes in the learning rate: their *direction* (up/down).
- We reject the Pearce-Hall Prediction-Error Learning: learning rate changes are insensitive to the surprise ($|\epsilon_t|$).
- We find convincing evidence that the learning rate is modulated by the sign of the prediction error autocovariance ($\epsilon_t \epsilon_{t-1}$)...
- ... but discover that this modulation is mostly observed when surprise is above the median level the agents experience individually.

2. Nature of the Data

Experiments where forecasts are for a self-referential "price system."

+ **Positive feedback**: forecasts of several forecasters are averaged and positively influence prices. E.g., High price expectations leads more purchases and supply did not adjust fast enough, leading to even higher prices.
– **Negative feedback**: averaged forecasts generate opposite reaction. E.g., High price expectations leads to over-production and hence price crashes.

- 6-9 participants per markets to ensure self-referentiality. (6-12 major players in the real world!)

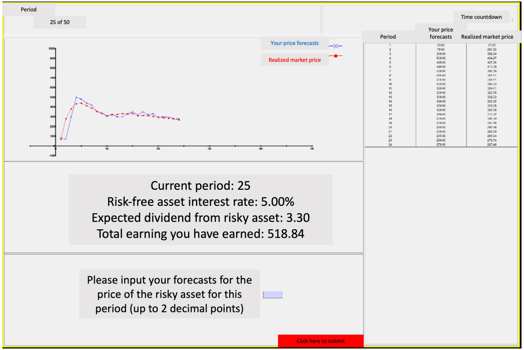

3. Results I: Imposing a constant learning rate α as in Bordalo ea

$$p^e_{t+1} = p^e_t + \bar{\alpha}(p_t - p^e_t)$$

Estimated αs should be higher than observed in Experiment 1-11, and lower than observed in 12-18??
Empirical signals (on top of model derivative) that constant learning rate α is not optimal in this game!

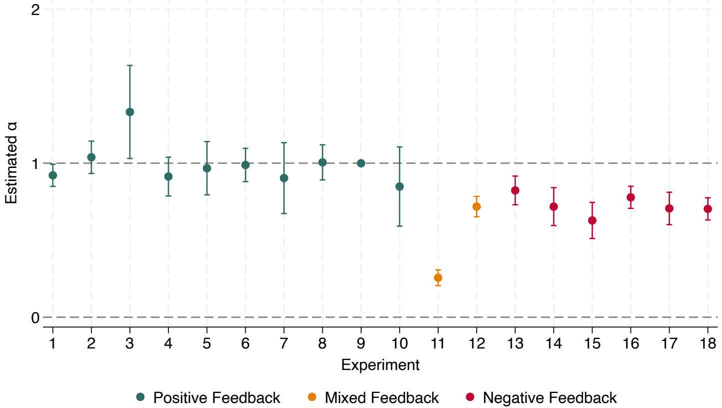

(Experiments 11+12: Mixed positive/negative feedback. Random-effects OLS model with cluster-robust standard errors for panels nested within participant level

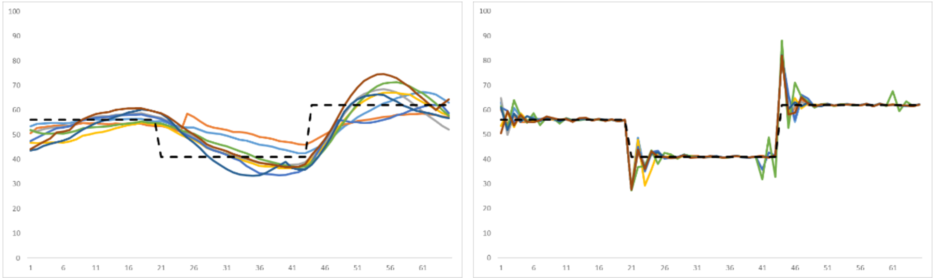

Sample price dynamics in positive (left panel) and negative feedback (right panel), representing data in Experiment 1-3, and Experiment 13-15, respectively, from Bao et al. (2012). The dashed line shows the equilibrium price under rational expectations.

4. Interim: Estimation

1. Direction of change:

$$Y_t = \begin{cases} +1 & \text{if } \frac{p^e_{t+1} - p^e_t}{\epsilon_t} - \frac{p^e_t - p^e_{t-1}}{\epsilon_{t-1}} > 0 \\ 0 & \text{if } \frac{p^e_{t+1} - p^e_t}{\epsilon_t} - \frac{p^e_t - p^e_{t-1}}{\epsilon_{t-1}} = 0 \\ -1 & \text{if } \frac{p^e_{t+1} - p^e_t}{\epsilon_t} - \frac{p^e_t - p^e_{t-1}}{\epsilon_{t-1}} < 0 \end{cases}$$

Need *ordered logit* model. Random effects. Cluster SE at market level. Bonferroni Correction of 54 estimations.

4. Interim: Estimation

1. Direction of change:

$$Y_t = \begin{cases} +1 & \text{if } \frac{p^e_{t+1} - p^e_t}{\epsilon_t} - \frac{p^e_t - p^e_{t-1}}{\epsilon_{t-1}} > 0 \\ 0 & \text{if } \frac{p^e_{t+1} - p^e_t}{\epsilon_t} - \frac{p^e_t - p^e_{t-1}}{\epsilon_{t-1}} = 0 \\ -1 & \text{if } \frac{p^e_{t+1} - p^e_t}{\epsilon_t} - \frac{p^e_t - p^e_{t-1}}{\epsilon_{t-1}} < 0 \end{cases}$$

   Need *ordered logit* model. Random effects. Cluster SE at market level. Bonferroni Correction of 54 estimations.

2. PH: Y^i_{t+1} closer to 1, when surprises $|\varepsilon^i_{t+1}| - |\varepsilon^i_t| > 0$

4. Interim: Estimation

1. Direction of change:

$$Y_t = \begin{cases} +1 & \text{if } \frac{p^e_{t+1} - p^e_t}{\epsilon_t} - \frac{p^e_t - p^e_{t-1}}{\epsilon_{t-1}} > 0 \\ 0 & \text{if } \frac{p^e_{t+1} - p^e_t}{\epsilon_t} - \frac{p^e_t - p^e_{t-1}}{\epsilon_{t-1}} = 0 \\ -1 & \text{if } \frac{p^e_{t+1} - p^e_t}{\epsilon_t} - \frac{p^e_t - p^e_{t-1}}{\epsilon_{t-1}} < 0 \end{cases}$$

   Need *ordered logit* model. Random effects. Cluster SE at market level. Bonferroni Correction of 54 estimations.

2. PH: Y^i_{t+1} closer to 1, when surprises $|\varepsilon^i_{t+1}| - |\varepsilon^i_t| > 0$
3. REA:

$$z^i_{t+1} = \gamma 1_{\{\epsilon^i_{t+1}\epsilon^i_t \geq 0\}} + \beta 1_{\{|\epsilon^i_{t+1}| < \text{median}(|\epsilon^i|)\}} + \delta 1_{\{|\epsilon^i_{t+1}| < \text{median}(|\epsilon^i|)\}} 1_{\{\epsilon^i_{t+1}\epsilon^i_t \geq 0\}})$$

   Such that:

$$Y^i_{t+1} = \begin{cases} -1 & \text{if } z^i_{t+1} + \eta^i_{t+1} < \mu^i_1, \\ 0 & \text{if } \mu^i_1 \leq z^i_{t+1} + \eta^i_{t+1} < \mu^i_2, \\ +1 & \text{if } \mu^i_2 \leq z^i_{t+1} + \eta^i_{t+1}, \end{cases}$$

5. Results II: (Direction of change in α) Pearce-Hall is rejected

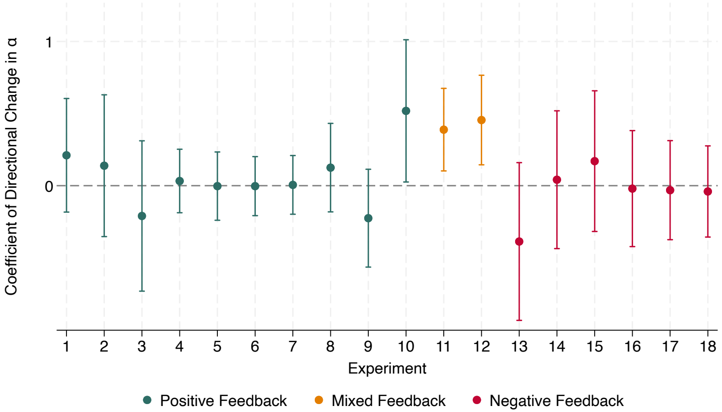

Logit model with $Y_t = 0$ added to outcome $Y_t = +1$ since the extra state is insignificant. Explanatory variable: binary, tracking whether surprise remained equal or increased. Random effects at the individual level. 95% CIs with Bonferroni correction.

5. Results III: (Direction of Change in α) Ordered Logit Random Coefficients Model Fully Supports MRAC

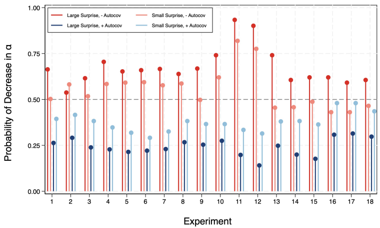

Decrease in α

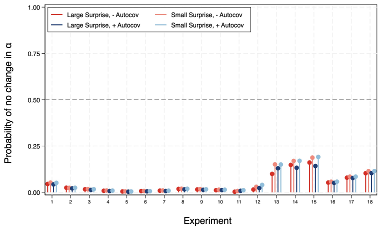

No Change in α

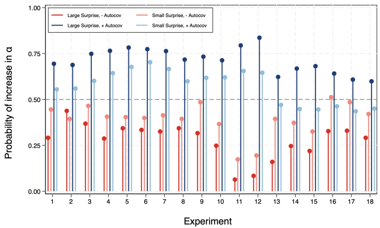

Increase in α

"Small" error = below individual's median prediction error size.

Conclusion

- In a setting where self-referentiality causes non-stationarity, participants adjust the learning rate in adaptive expectations as a function of auto-covariance of prediction errors, provided surprise (absolute value of prediction error) is sufficiently high.
- This confirms that subjects are implementing robust control in the form of MRAC (Model Reference Based Control).
- The findings are consistent with behavior and neural processing in other high-cognition tasks, as well as in much simpler sensorimotor tasks.

Thank you !

Jiaoying Pei: jp2127@cam.ac.uk

References

Anufriev, M. and C. Hommes (2012). Evolutionary selection of individual expectations and aggregate outcomes in asset pricing experiments. *American Economic Journal: Microeconomics 4*(4), 35–64.

Bao, T., J. Duffy, and C. Hommes (2013). Learning, forecasting and optimizing: An experimental study. *European Economic Review 61*, 186–204.

Bao, T., S. Füllbrunn, J. Pei, and J. Zong (2024). Reading the market? expectation coordination and theory of mind. *Journal of Economic Behavior & Organization 219*, 510–527.

Bao, T. and C. Hommes (2019). When speculators meet suppliers: Positive versus negative feedback in experimental housing markets. *Journal of Economic Dynamics and Control 107*, 103730.

Bao, T., C. Hommes, and T. Makarewicz (2017). Bubble formation and (in) efficient markets in learning-to-forecast and optimise experiments. *The Economic Journal 127*(605), F581–F609.

Bao, T., C. Hommes, J. Sonnemans, and J. Tuinstra (2012). Individual expectations, limited rationality and aggregate outcomes. *Journal of Economic Dynamics and Control 36*(8), 1101–1120.

Bordalo, P., N. Gennaioli, Y. Ma, and A. Shleifer (2020). Overreaction in macroeconomic expectations. *American Economic Review 110*(9), 2748–2782.

Bossaerts, P. and P. Hillion (1999). Implementing statistical criteria to select return forecasting models: what do we learn? *Review of Financial Studies 12*(2), 405–428.

Carvalho, C., S. Eusepi, E. Moench, and B. Preston (2023). Anchored inflation expectations. *American Economic Journal: Macroeconomics 15*(1), 1–47.

Csáji, B. C. and L. Monostori (2008). Adaptive stochastic resource control: a machine learning approach. *Journal of Artificial Intelligence Research 32*, 453–486.

d'Acremont, M. and P. Bossaerts (2016). Neural mechanisms behind identification of leptokurtic noise and adaptive behavioral response. *Cerebral Cortex 26*(4), 1818–1830.

Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica 57*(2), 357–384.

Hansen, L. P. and T. J. Sargent (2001). Acknowledging misspecification in macroeconomic theory. *Review of Economic Dynamics 4*(3), 519–535.

Marcet, A. and T. J. Sargent (1989). Convergence of least squares learning mechanisms in self-referential linear stochastic models. *Journal of Economic theory 48*(2), 337–368.

Pearce, J. M. and G. Hall (1980). A model for pavlovian learning: variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological review 87*(6), 532.

Simon, H. A. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics*, 99–118.

Sutton, R. S. (1992). Adapting bias by gradient descent: An incremental version of delta-bar-delta. In *AAAI*, Volume 92, pp. 171–176. Citeseer.