

Jump, Robert

Working Paper

Results on the Stability of a Simple Wage Posting Model

School of Economics Discussion Papers, No. 1319

Provided in Cooperation with:

University of Kent, School of Economics

Suggested Citation: Jump, Robert (2013) : Results on the Stability of a Simple Wage Posting Model, School of Economics Discussion Papers, No. 1319, University of Kent, School of Economics, Canterbury

This Version is available at:

<https://hdl.handle.net/10419/105707>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

University of Kent
School of Economics Discussion Papers

**Results on the Stability of a Simple
Wage Posting Model**

Robert Jump

November 2013

KDPE 1319



Results on the Stability of a Simple Wage Posting Model*

Robert Jump[†]

School of Economics, University of Kent

November 2013

Abstract

This paper presents results on the stability of the wage dispersion model presented in Mortensen (2003). Specifically, we test four “positive definite” learning processes on a single parameterisation of the underlying model, and submit the most successful to a thorough sensitivity analysis. The general result of existing studies of the stability of price dispersion models is that learning processes can converge on limiting distributions that qualitatively match the equilibrium distribution. In contrast, the most successful process considered in this paper can converge on a limiting distribution that quantitatively matches the equilibrium distribution.

Keywords: Price dispersion, Search market equilibrium, Reinforcement learning.

JEL Classification: C62, C63, D83, J31.

*Funding for this research was provided by the ESRC (grant number: ES/J500148/1). I would like to thank Mathan Satchi for his ongoing help, and Jagjit Chadha for useful comments on an earlier draft. All remaining errors and omissions are the responsibility of the author.

[†]School of Economics, University Of Kent, Canterbury CT2 7NP, UK: rgj4@kent.ac.uk, +44 (0)1227 827952.

1 Introduction

This paper presents results on the stability of the simple wage posting model presented in Mortensen (2003: 16 - 20), which is a simplification of Burdett and Mortensen (1998). The model consists of a single period, in which m firms face n households. Each firm posts a single wage offer to a randomly chosen household, and each household accepts the highest offer received. The number of offers received by any given household is then binomially distributed, with “probability of success” equal to $1/n$, and “sample size” equal to m . This environment, in which households cannot observe every wage offer on the market, but are not confined to taking or leaving a single offer, means that firms have to trade off the benefits of a higher wage offer (higher probability of acceptance) against those of a lower wage offer (higher profit given acceptance). These opposing forces, in turn, give rise to a non-degenerate wage offer distribution in equilibrium. Assuming a large market, such that the binomial distribution can be approximated by the Poisson distribution, the probability of acceptance is as follows:

$$P(F(w), \lambda) = \sum_{x=0}^{\infty} \frac{e^{-\lambda} \lambda^x}{x!} F(w)^x = e^{-\lambda[1-F(w)]} \sum_{x=0}^{\infty} \frac{e^{-\lambda F(w)} (\lambda F(w))^x}{x!} = e^{-\lambda[1-F(w)]} \quad (1)$$

Here, $\lambda = m/n$ denotes market tightness, and $F(w)$ the wage offer distribution. Expected profit is then given by the difference between the marginal revenue product p and the wage offer w given the probability of acceptance. If we suppose that there is a common reservation wage b , the equilibrium offer distribution can be found by appealing to an equal profit condition. As the reservation wage will only be accepted if it is the only offer received, the expected profit of offering the reservation wage is independent of the offer distribution, and determines the expected profit of all other offers:

$$\pi(p, w, 0) = (p - b)e^{-\lambda} = (p - w)e^{-\lambda[1-F(w)]} = \pi(p, w, F(w)) \quad (2)$$

Solving eq 2 for F then yields the unique offer c.d.f., which can be solved for $F = 1$ to yield the upper support of the distribution, \bar{w} :

$$F(w) = \frac{1}{\lambda} \ln \left(\frac{p - b}{p - w} \right) \quad (3)$$

$$\bar{w} = p - e^{-\lambda}(p - b) \quad (4)$$

Thus the offer distribution is non-degenerate, with upper support less than the marginal revenue product of a match¹. Although this model is, in a structural sense, extremely simple, the situation it describes is rather complex from the point of view of the individual firm. As

¹See Mortensen (2003: 16 - 20) for the proof that the distribution described by eqs 3 and 4 is the unique equilibrium.

a lower wage offer trades off a higher profit given acceptance against a lower probability of acceptance, an individual firm has to know the probability of acceptance, conditional on all wage offers, in order to make an informed decision in regards to its individual offer *ex ante*. Unless one assumes that this knowledge exists *ex ante*, it is not at all clear how the individual firm should behave. As a result, the existing literature on the stability of price dispersion models utilises general learning processes borrowed from the machine learning literature, in particular, replicator dynamics and reinforcement learning. In terms of general results, Hopkins and Seymour (1996) demonstrates that a class of “positive definite” dynamics, in which a large economy assumption allows noise to be ignored, can result in the stability of dispersed price equilibria under learning. A more recent paper explores the implications of the large economy (or large population) assumption and the ability of finite population learning mechanisms to match deterministic processes (Gao et al 2013).

Although these results are encouraging, they are not reflected in the literature examining the ability of specific learning processes to match dispersed price equilibria. As with Hopkins and Seymour (1996), Cason et al (2005) examines the stability of the price dispersion model presented in Burdett and Judd (1983). Four different replicator dynamics are studied. The limiting distributions to which the majority of processes tend are found to be qualitatively, rather than quantitatively, similar to the equilibrium distribution. In addition, one process does not converge at all, but tends to a limit cycle around the equilibrium. Waldeck and Darimon (2006) conducts a fully specified microsimulation of a market in which individual firms independently adapt via reinforcement learning, in the context of the price dispersion model of Varian (1980). They consider a model with 1000 buyers and 20 sellers, where each seller posts a price from a set of size 100. The computational cost of such a model is relatively high, so that the authors only consider one learning process. As with Cason et al (2005), they conclude that the learning process does not converge to the equilibrium, but to a qualitatively similar limiting distribution.

The foregoing suggests that the stability of price dispersion models is sensitive to the exact learning process employed. As such, it seems prudent to examine a number of different processes, which largely precludes full microsimulation, as with Waldeck and Darimon (2006)². The approach taken in this paper, therefore, is to study a number of “positive definite” processes, as with Hopkins and Seymour (1996) and Cason et al (2005). Section 2 presents four different processes, which are subjected to an initial test in section 3 for a single fundamental parameterisation of the underlying model. The most successful process is then subjected to a sensitivity analysis in section 4. The conclusion reached is that this process can converge on a limiting distribution that quantitatively, rather than qualitatively, matches the equilibrium distribution. This is a unique result in the literature. This result is not sensitive to the fundamental parameterisation, but is sensitive to the parameterisation of the learning process. Unfortunately, for certain parameterisations the process becomes extremely unstable, which qualifies the results somewhat. Nevertheless, the conclusion is encouraging for the stability of this class of models.

2 Candidate Learning Processes

Consider a large number of firms, facing a large number of households, such that the ratio of firms to households is λ . The environment is analogous to that described by the Mortensen

²For this reason other approaches, such as genetic algorithms (e.g. McCarthy 2009), are also rejected.

(2003) wage dispersion model: each firm posts a wage offer to a randomly picked household, and each household subsequently accepts the highest offer received (if any). The marginal revenue product of a match is p , and the reservation wage b is public knowledge. The set of possible wage offers is an equally spaced grid: $S = \{w_1, \dots, w_i, \dots, w_l\}$, with $w_{i+1} > w_i$, $w_1 = b$, and $w_l = p$. Given the above, assume that the number of firms and households is much larger than l , such that there is, at any point in time, a large number of firms playing each strategy. The time t average profit of each strategy is then given by the following:

$$\pi_{it}(p, w_i, \tilde{F}_t) = (p - w_i)e^{-\lambda[1 - \tilde{F}_t(i)]} \quad (5)$$

Here, \tilde{F}_t denotes the cumulative distribution of firms across the strategy set at time t , which may or may not be different from the equilibrium distribution F . Likewise, $\tilde{F}_t(i)$ denotes the value of \tilde{F}_t evaluated at strategy i , ie the proportion of firms playing strategies 1 to i inclusive at time t . The processes that are examined here borrow from the reinforcement learning literature in supposing that each strategy has a fitness measure associated with it, and that these fitness measures determine the proportion of firms playing each strategy. There are then two separate problems in constructing learning processes: first, how the densities of firms playing each strategy are determined, given the fitness measures, and second, how those fitness measures are determined.

Learning Process:

Suppose we have a function A to determine the densities of firms playing each of the l strategies, and a function B to determine those strategies' fitness measures. The general learning process can then be described by the following pseudo-code:

1. $\tilde{f}_{it} = A_a[\phi_{1t}, \dots, \phi_{it}, \dots, \phi_{lt}]$
2. $\tilde{F}_t \leftarrow \tilde{f}_t$
3. $\pi_{it} = (p - w_i)e^{-\lambda[1 - \tilde{F}_t(i)]}$
4. $\phi_{it+1} = B_b[\phi_{1t}, \dots, \phi_{it}, \dots, \phi_{lt}, \pi_{1t}, \dots, \pi_{it}, \dots, \pi_{lt}]$

As above, \tilde{F}_t denotes the cumulative distribution of firms over strategies, and \tilde{f}_{it} denotes the density of firms playing the i th strategy, at time t . This is determined by A , which takes fitness measures as arguments, denoted ϕ_{it} . These fitness measures comprise the model's l state variables, and thus we require an initial distribution of fitness levels to fully specify the model. Given these values, which determine \tilde{F}_1 , intra-period expected profits are updated for each strategy, which allows the fitness measures to be updated. This is determined by B , which allows \tilde{f}_{t+1} to be calculated. Hence, the process can be iterated, given initial conditions, to examine its limiting distribution.

Strategy Selection Functions:

Two possible functional forms for A are considered:

$$A_1: \tilde{f}_{it} = \frac{\phi_{it}}{\sum_{i=1}^l \phi_{it}}$$

$$A_2: \tilde{f}_t = \frac{e^{\phi_{it}/\tau}}{\sum_{i=1}^l e^{\phi_{it}/\tau}}$$

Denoting the space of vectors $[\phi_{1t}, \dots, \phi_{it}, \dots, \phi_{lt}]$ as Φ , and by Z the space of vectors $[\tilde{f}_{1t}, \dots, \tilde{f}_{it}, \dots, \tilde{f}_{lt}]$, A_1 and A_2 are mappings from $\Phi \rightarrow Z$. With A_1 , in any given period, the proportion of agents playing strategy i is equal to the relative fitness of that strategy. This is straightforward, and follows the usual manner in which replicator dynamics and reinforcement learning determine the proportion of agents over strategies. With A_2 , the proportion of agents playing strategy i is exponentially related to strategy i 's relative fitness; that is, a strategy i that is twice as fit as a strategy j is played by more than twice the number of agents. The extent to which this is the case is determined by the intensity parameter τ , such that a lower τ increases the rate at which agents choose relatively profitable strategies. This is variously known as ‘‘softmax’’ selection (Sutton and Barto 1998: 30) or ‘‘logit dynamics’’ (Cason et al *op. cit.*).

Fitness Updating Functions:

Two possible functional forms for B are considered:

$$B_1: \phi_{it+1} = \begin{cases} \phi_{it} + \alpha(\pi_{it} - \phi_{it}) & \text{if } \phi_{it} + \alpha(\pi_{it} - \phi_{it}) > 0 \\ 0 & \text{if } \phi_{it} + \alpha(\pi_{it} - \phi_{it}) \leq 0 \end{cases}$$

$$B_2: \phi_{it+1} = \begin{cases} \phi_{it} + \alpha(\pi_{it} - \bar{\pi}_t) & \text{if } \phi_{it} + \alpha(\pi_{it} - \bar{\pi}_t) > 0 \\ 0 & \text{if } \phi_{it} + \alpha(\pi_{it} - \bar{\pi}_t) \leq 0 \end{cases}$$

Denoting the space of vectors $[\pi_{1t}, \dots, \pi_{it}, \dots, \pi_{lt}]$ as Π , B_1 and B_2 are mappings from $\Phi \times \Pi \rightarrow \Phi$. With B_1 , fitness levels are updated as per the standard reinforcement learning rule, where the fitness measure of strategy i in any period is an exponentially weighted moving average of past profitability. B_2 is known as a ‘‘reinforcement comparison’’ algorithm in the machine learning literature (Sutton and Barto 1998: 41), where the fitness measure of strategy i is updated by comparing that strategy’s intra-period profit to the intra-period arithmetic average of all strategies’ profits, $\bar{\pi}_t$. As we are considering an aggregated system, this becomes a type of social learning, and similar in spirit to replicator dynamics. The foregoing, by the different possible combinations of updating functions, gives four separate learning processes, which will be referred to as processes A_1B_1 , A_1B_2 , A_2B_1 , and A_2B_2 . We are interested primarily in the similarity between the limiting distributions to which these processes tend and the equilibrium distribution described by eqs 3 and 4. Section 3 compares the convergence results of the four learning processes for a single parameterisation




Figure 1: A_1B_1 , $\alpha = 0.1$




Figure 2: A_1B_1 , $\alpha = 0.5$




Figure 3: A_1B_1 , $\alpha = 0.1$




Figure 4: A_1B_1 , $\alpha = 0.5$

of the underlying model. In general, their performance is as expected, given the results of the existing literature. However, A_2B_1 is found to perform extremely well for certain parameterisations, and this is the process that is subjected to a further sensitivity analysis in section 4.

3 Results: Comparative Analysis

The four learning processes defined above are iterated numerically to generate their limiting distributions (using Matlab code available upon request). In order to measure both the similarity between the limiting distributions and the equilibrium distribution, and the speed at which the processes converge on their respective limiting distributions, we use the Kolmogorov-Smirnov (KS) statistic:

$$KS_t = \sup_i |F(i) - \tilde{F}_t(i)| \quad (6)$$

The KS statistic is the supremum of the set of absolute differences between F and \tilde{F}_t in any given period, and the KS statistic at a process's limiting distribution is the greatest absolute difference between that distribution and the equilibrium distribution³.

³Alternative distance metrics could have been used here; the KS statistic is chosen to ensure comparability with the relevant literature, e.g. Waldeck and Darimon (2006).




Figure 5: A_1B_2 , $\alpha = 0.1$




Figure 6: A_1B_2 , $\alpha = 0.5$




Figure 7: A_1B_2 , $\alpha = 0.1$




Figure 8: A_1B_2 , $\alpha = 0.5$

Thus *qualitatively match* is precisely defined only in a relative sense; that is, a process with a lower KS statistic at its limiting distribution is more similar to the equilibrium distribution than a process with a higher KS statistic. In turn, as the KS statistic approaches zero, the limiting distribution approaches the equilibrium distribution. Sections 3.1 - 3.4 describe the convergence results of the four learning processes, with the fundamental parameterisation as follows: $\lambda = 1$, $b = 1$, and $p = 2$. Although we consider different learning parameters, the size of the strategy set $l = 1000$ throughout, and initial fitness measures are uniform for each strategy.

3.1 A_1B_1

Figures 1 - 4 illustrate the convergence properties of A_1B_1 . Figures 1 and 2 show the limiting distribution of the learning process (dashed line) against the equilibrium distribution (solid line) for $\alpha = 0.1$ and $\alpha = 0.5$ respectively. Figures 3 and 4 show the evolution of the KS statistic over iterations 1 - 1000 for the same parameter values. It is instantly apparent, from visual inspection of figures 1 and 2, that the limiting distribution of A_1B_1 is not particularly close to the equilibrium distribution, and does a particularly bad job of estimating \bar{w} , with no wage offers disappearing from the strategy set at all. Furthermore, the limiting distribution is invariant to α , the sole effect of which is to determine the adjustment speed, and the KS statistic at $t = 1000$ is approximately 0.21. Although A_1B_1 is very stable around its

limiting distribution, therefore, it is not particularly successful at learning the equilibrium distribution.

3.2 A_1B_2

Figures 5 - 8 illustrate the convergence properties of A_1B_2 , for the same two parameterisations as above. As with A_1B_1 , the limiting distribution is invariant to the choice of α , which governs the adjustment speed. With this learning process, however, the limiting distribution is much closer to the equilibrium distribution, with the KS statistic at $t = 1000$ approximately 0.085. Moreover, the process does a relatively good job of estimating \bar{w} , with the majority of wage offers above \bar{w} disappearing from the limiting distribution. However, where A_1B_1 smoothly approaches its limiting distribution, A_1B_2 exhibits a degree of cyclical adjustment in the KS statistic. This is indicated by the KS statistic rapidly falling to 0.1 in each parameterisation, after jumping up again and steadily approaching 0.085 over the remaining iterations. Correspondingly, A_1B_2 takes a significantly greater number of iterations to reach its minimum KS statistic than A_1B_1 . Particularly, where A_1B_1 with $\alpha = 0.5$ exhibits extremely rapid convergence, A_1B_2 takes approximately 600 iterations. Despite this, we can conclude that A_1B_2 does a relatively good job of learning the equilibrium distribution, with particular success at estimating the upper support. This result is in line with the existing literature referred to above.

3.3 A_2B_1

Unlike A_1B_1 and A_1B_2 , A_2B_1 has two parameters governing the learning process: the adjustment parameter α as before, and the strategy choice parameter τ . Figures 9 - 12 show the limiting distribution of the learning process for four parameter combinations: $\{\alpha = 0.1, \tau = 0.5\}$, $\{\alpha = 0.1, \tau = 0.05\}$, $\{\alpha = 0.005, \tau = 0.05\}$, and $\{\alpha = 0.005, \tau = 0.005\}$. As before, α does not affect the limiting distribution. However, the strategy choice parameter does have an effect on the limiting distribution, and a judicious choice can greatly improve the ability of this learning process to match the equilibrium distribution. In general, it is the case that the KS statistic decreases with τ ; that is, the ability of A_2B_1 to learn the equilibrium distribution improves as τ decreases. This is an interesting result, as the rate of flow of agents from low to high fitness strategies increases as τ decreases. Intuitively, it is this extra “degree of freedom” which improves the performance of this process relative to A_1B_1 and A_1B_2 . Unfortunately, as τ passes a certain threshold, the process fails to converge at all, and can display extremely chaotic behaviour. An example is given in figure 13, which plots the KS statistic over iterations 1 - 150 for $\{\alpha = 0.5, \tau = 0.02\}$. Furthermore, the threshold value of τ at which the process becomes unstable is dependent on α , and figure 14 provides a stability plot for different combinations of α and τ , from which the relatively large region of unstable parameterisations is immediately apparent⁴. Despite this, the success of the *stable* parameterisations of A_2B_1 is unambiguously greater than the processes examined thus far. In fact, jointly decreasing α and τ results in an arbitrarily low KS statistic - the parameterisation illustrated in figure 12, for example, has a KS statistic

⁴In figure 14, instability is defined as an absolute difference of 0.001 or more between the KS statistics at iterations 999 and 1000.




Figure 9: A_2B_1 , $\alpha = 0.1$, $\tau = 0.5$




Figure 10: A_2B_1 , $\alpha = 0.1$, $\tau = 0.05$




Figure 11: A_2B_1 , $\alpha = 0.005$, $\tau = 0.05$




Figure 12: A_2B_1 , $\alpha = 0.005$, $\tau = 0.005$




Figure 13: A_2B_1 , $\alpha = 0.5$, $\tau = 0.02$




Figure 14: A_2B_1 Stability Plot




Figure 15: A_2B_2 , $\alpha = 0.1$, $\tau = 0.3$




Figure 16: A_2B_2 , $\alpha = 0.5$, $\tau = 0.5$

of 0.0115 at its limiting distribution. At this point, therefore, we can tentatively conclude that A_2B_1 has the ability to quantitatively match the equilibrium distribution described by eqs 3 and 4.

3.4 A_2B_2

The final positive definite learning process that we consider incorporates “logit dynamics” into the “reinforcement comparison” fitness updating algorithm. As incorporating the former into the basic fitness updating algorithm significantly improved its performance, it might be imagined that A_2B_2 would emerge as the most successful process. Unfortunately, the tendency of A_2 towards instability appears to interact with the cyclical adjustment of B_2 in such a way that A_2B_2 does not converge to a limiting distribution for any combination of parameter values. Instead, the process produces explosive cyclical motion in the KS statistic. Figures 15 and 16 illustrate this for the parameterisations $\{\alpha = 0.1, \tau = 0.3\}$ and $\{\alpha = 0.5, \tau = 0.5\}$, respectively. Finally, therefore, we can conclude that A_2B_2 is not successful as a learning process in the context of the wage dispersion model considered here.

3.5 Discussion

To summarise, A_1B_1 is stable around its limiting distribution, but that distribution is relatively dissimilar to the equilibrium distribution, with a KS statistic of 0.21. At the other extreme, A_2B_2 produces explosive cyclical motion for all parameterisations, and thus fails completely as a learning process. The two relatively successful processes are A_1B_2 and A_2B_1 . The former converges reliably, albeit slowly, with a KS statistic of 0.085 at its limiting distribution. In comparison, A_2B_1 can achieve an arbitrarily low KS statistic by reducing α and τ jointly. This is a delicate process, however, as after a certain threshold any reduction in τ with α fixed causes extreme instability. Despite this, the apparent match between the limiting distribution of A_2B_1 and the equilibrium is a unique result in the literature. As such, this is the process that is subjected to a more rigorous sensitivity analysis in section 4. Particularly, we test its ability to learn the equilibrium for four different fundamental parameterisations, and the region of unstable $\{\alpha, \tau\}$ combinations for those parameterisations.

4 Results: A_2B_1 Sensitivity Analysis

As demonstrated above, the KS statistic corresponding to the limit distribution of A_2B_1 decreases with τ . Given this, the threshold value of τ at which the process becomes unstable increases with α ; hence for the fundamental parameterisation considered above, the KS statistic is minimised in the most south-westerly corner of the stable region of figure 14. When considering the sensitivity of the performance of this learning process, therefore, the primary question of interest is whether the region of unstable $\{\alpha, \tau\}$ combinations significantly differs for different fundamental parameterisations. The four parameterisations of the underlying model that A_2B_1 are tested against are as follows:

P1: $\lambda = 0.5, p = 2, b = 1$

P2: $\lambda = 1.5, p = 2, b = 1$

P3: $\lambda = 1, p = 3, b = 1$

P4: $\lambda = 1, p = 2, b = 1.5$

That is, in comparison to the parameterisation considered in section 3, P1 decreases λ , P2 increases λ , P3 increases $(p - b)$, and P4 decreases $(p - b)$. Figures 17 - 20 present stability plots for each of these parameterisations, calculated in the same way as figure 14, above. As can be seen in figures 17 and 18, varying λ with p and b fixed does have an effect on the region of unstable $\{\alpha, \tau\}$ combinations, but the effect is not particularly pronounced. In contrast, varying $(p - b)$ with λ fixed has a significant effect; reducing the difference between the marginal revenue product and the reservation wage decreases the unstable region substantially, whilst increasing this difference enlarges the region. It is not clear what the economic intuition for this result is, although it is worth noting that the instability generated still corresponds to cyclical motion in the KS statistic, as in figure 13, so that the process does not halt at any point.

The foregoing indicates that the success of the learning process is relatively unaffected by the choice of λ , and rather more sensitive to the choice of p and b . However, for an appropriate choice of α and τ , the limiting distribution to which A_2B_1 converges can still achieve an arbitrarily low KS statistic for both P3 and P4. This is illustrated in figures 21 and 22, which compare the limiting distribution and equilibrium distribution for P3 and P4, with $\{\alpha = 0.005, \tau = 0.01\}$ and $\{\alpha = 0.003, \tau = 0.003\}$, respectively. The KS statistic for the former is 0.0154, and the KS statistic for the latter is 0.0146. Finally, although we do not present the results here, the process is largely unaffected by the size of the strategy set. Reducing l substantially (ie < 100) does affect the stability properties, but mainly in the time taken to convergence rather than the fact of convergence itself. Increasing l past 1000, on the other hand, has no material effect on the results. Similarly, randomising the initial distribution of fitness measures, rather than specifying uniformity, has no effect on the limiting distribution.

The conclusions of this section are, it is fair to say, mixed. Not only is the most successful process considered in section 3 unstable over a relatively large parameter region, but this region itself is affected by the fundamental parameterisation of the underlying model. An important consequence of this is that, for a given parameterisation of A_2B_1 , this learning




Figure 17: P1 Stability Plot




Figure 18: P2 Stability Plot




Figure 19: P3 Stability Plot




Figure 20: P4 Stability Plot




Figure 21: P3, $\alpha = 0.005$, $\tau = 0.01$




Figure 22: P4, $\alpha = 0.003$, $\tau = 0.003$

process will *not* dominate the alternative processes examined in section 3 for any random fundamental parameterisation. Thus there is an issue, so far unexamined here, concerning learning under parameter uncertainty. Despite this, it remains the case that for a judicious $\{\alpha, \tau\}$ combination, A_2B_1 can achieve an arbitrarily low KS statistic at its limiting distribution for a variety of fundamental parameterisations. As such, the justification remains for concluding that the most successful process considered here can converge on a limiting distribution that quantitatively matches the equilibrium distribution of the Mortensen’s simple version of the Burdett and Mortensen (1998) wage dispersion model. This is an encouraging result for this class of models, and suggests that the exercise might be profitably repeated for different price dispersion models. At the same time, variants of A_2B_1 might be found that can achieve its convergence success without the associated problem of instability.

5 Concluding Remarks

The four candidate learning processes considered in this paper vary widely in their ability to learn the equilibrium distribution of Mortensen’s simple wage posting model. The two basic processes, A_1B_1 and A_1B_2 , are in line with results reached by Cason et al (2005) and Waldeck and Darimon (2006); that is, convergence to a limiting distribution qualitatively similar to the equilibrium distribution. The least successful process, A_2B_2 , generates explosive cyclical motion for every parameterisation, which is a failure that has not been reported in the existing literature. The most successful process, on the other hand, can converge on a limiting distribution that can quantitatively match the equilibrium distribution. The word “can” should be emphasised here, as the similarity between the limiting distribution of A_2B_1 and the equilibrium distribution is sensitive to the learning parameters. Moreover, although simultaneously decreasing α and τ improves learning success, there is a threshold ratio of the two parameters at which the process generates severe instability. This problem is magnified by the ratio being sensitive to the fundamental parameterisation of the underlying model, which leads to a significant problem concerning learning under parameter uncertainty.

Despite these qualifications, it remains the case that a judicious choice of α and τ allows A_2B_1 to learn the equilibrium distribution to an arbitrarily high degree of accuracy regardless of the fundamental parameterisation. This is a unique result in the literature, and is encouraging for further work in this field. Two avenues immediately suggest themselves. First, the exercise might be profitably repeated for different models in this class - so far, only three models have been examined (the model examined here, and the models examined in Cason et al (2005) and Waldeck and Darimon (2006)). Second, variants on A_2B_1 might be found which can achieve similar learning success whilst avoiding the problem of instability. This is necessary to deal with the issue of parameter uncertainty, and a possible line of attack would be real-time control of τ . Particularly, if τ were permitted to react to volatility in the distance metric, a learning process that ensures convergence regardless of the fundamental parameterisation might be found.

References

- [1] Burdett, K., and Judd, K. 1983. Equilibrium Price Dispersion. *Econometrica*. **51**, 955-969.
- [2] Burdett, K., and Mortensen, D. 1998. Wage Differentials, Employer Size, And Unemployment. *International Economic Review*. **39**, 257-273.
- [3] Cason, T., Friedman, D., and Wagner, F. 2005. The Dynamics Of Price Dispersion, Or Edgeworth Variations. *Journal Of Economic Dynamics And Control*. **29**, 801-822.
- [4] Gao, X., Zhong, W., and Mei, S. 2013. Stochastic Evolutionary Game Dynamics And Their Selection Mechanisms. *Computational Economics*. **41**, 233-247.
- [5] Hopkins, E., and Seymour, R. 1996. Price Dispersion: An Evolutionary Approach. *Discussion Paper Series - University Of Edinburgh Department Of Economics*.
- [6] McCarthy, I. 2009. Simulating Sequential Search Models With Genetic Algorithms: Analysis Of Price Ceilings, Taxes, Advertising And Welfare. *Computational Economics*. **34**, 217-241.
- [7] Mortensen, D. 2003. *Wage Dispersion: Why Are Similar Workers Paid Differently?* Cambridge MA: MIT Press.
- [8] Sutton, R., and Barto, A. 1998. *Reinforcement Learning: An Introduction*. Cambridge MA: MIT Press
- [9] Varian, H. 1980. A Model Of Sales. *American Economic Review*. **70**, 651-659.
- [10] Waldeck, R., and Darimon, E. 2006. Can Boundedly Rational Sellers Learn To Play Nash? *Journal Of Economic Interaction And Coordination*. **1**, 147-169.