An Agent-based Approach to the Gravity Model of Linguistic Diffusion

Xiaoxing Yu†

Project advisor: Christina J. Edholm‡

Abstract. Mathematical models of language change promise to provide insight on how community-wide changes emerge from individual speaker interactions. We develop an agent-based model to test the validity of the gravity theory of linguistic diffusion. Our model uses an exemplar model and the functional-teleological explanation of chain shifts to build up a picture of how individual agents change their vowel systems, and studies the emergent patterns across language communities. We find that the predictions of the gravity model hold true in a subset of cases determined at least in part by patterns of interaction and the particularities of how vowels are produced and perceived. To be specific, the gravity model's predictions will hold true in situations where two interacting communities have categories within each other's tolerance thresholds.

1. Introduction.

Communities of language users in social contact with each other often influence each other's language norms. Situations of extensive interborrowing of features at all linguistic levels are well-known and well-documented, as in the case of the languages of mainland Southeast Asia, which share similarities in the phonological, morphological, and syntactic structure[1] [1]. Another often cited example of languages in contact is that of the so-called "standard average European", which refers to the fact that a dozen or so prominent linguistic features are common to a collection of genealogically[2] unrelated or distantly related languages of Europe [4]. Investigating the mechanisms by which the languages spoken by geographically close communities in social contact influence one another is hence a crucial

†Pomona College, Claremont, CA (xyaa2021@mymail.pomona.edu)
‡Scripps College, Claremont, CA (cedholm@scrippscollege.edu)

[1] Roughly, the structure of their sounds, words, and sentences.

[2] Two languages are considered genealogically related if they can be shown to descend from a common ancestor.

component of understanding language change and complements genealogical, "within-speaker community" approaches.

Speaker communities are themselves hierarchical, since larger communities may contain smaller communities, and inhomogeneous, since not all speakers in a community speak in exactly the same way. The United States may be considered a large speaker community sharing a common, structurally upheld standard (Mainstream U.S. English, MUSE). However, it also comprises smaller communities at the city, town, or social network level. As the *Atlas of North American English* (ANAE) [5] documents, the United States is home to an incredibly diverse array of English varieties, subsets of which are in sociolinguistic contact. Many of the same techniques and theories that are used to understand the interaction of speaker communities of different languages can be used to understand the interaction of speaker communities of different varieties of a single language, usually referred to as "dialects" [8, ch. 1]. In fact, we expect linguistic features to spread more easily between two speaker communities that can understand each other's speech than between two communities that cannot.

Our work looks at exactly such a situation of interdialect contact. In particular, we build on the work of Stanford and Kenny in developing an agent-based model for the Northern Cities Shift (NCS), a sound change that has taken root in U.S. cities around the Great Lakes[7]. Although the NCS (for more detail see Appendix A in the supplementary material) is taken as a starting point for the model, we abstract away from the particularities of the NCS to arrive at a model that can be applied to other vowel chain shifts. In order to retain comparability to Stanford and Kenny's results and to not overcomplicate the model, we adopt a number of simplifying assumptions, many of which Stanford and Kenny employ themselves. For instance, we assume that vowels may be modeled one-dimensionally and ignore the mediating role of social networks. We view these simplifications as necessary to establish a first model that can reproduce qualitative predictions, and leave additional modifications to future research. For more detail on these simplifications, see § 2.

The present project's central research question is whether and how speaker-to-speaker interactions and changes in speaker exemplars alone are sufficient to reproduce the predictions of the gravity model. In particular, it remains to be verified whether increasing population density and decreasing distance lead to increased rates of diffusion, and whether a mechanistic model can reproduce the phenomenon of a change "skipping over" an intervening settlement. Section 3 will show that these predictions hold true when two interacting communities have categories within each other's tolerance thresholds and that it is thus necessary to also consider patterns of interaction and the particularities of vowel production and perception.

Section 2 describes the model according to the ODD protocol for agent-based models [3]. Section 3 shows and analyzes the results obtained from the model, and conducts sensitivity analysis on the model parameters. Finally, § 4 summarizes the results, concludes that the predictions of the gravity model hold true in a subset of cases, evaluates the model's limitations, and proposes future research directions.

2. Methods. Our model is agent-based and constructed in the NetLogo software[3]. Agent-based models (ABMs) start from a set of assumptions about how individual agents in a system behave and then simulate the evolution of the system as a whole. In this way, ABMs capture emergent, individual-level behaviors that deterministic models cannot. The model description follows the ODD (Overview, Design concepts, Details) protocol for describing individual- and agent-based models [3], as updated by [2]. Our model builds off the previous work by Stanford and Kenny [7]. For the full ODD, see Appendix B.

The theoretical foundation of our model is the gravity model of language contact, as described in [9] (for more on the gravity model, see Appendix A). The gravity model posits that interaction rates between two language communities correspond with the rate of linguistic diffusion between the communities. Crucially, this implies that it is possible for language features to spread from a large, densely populated community to a second densely populated

[3]Please reach out to xyaa2021@mymail.pomona.edu for the full code.

community while skipping over geographically intervening sparsely populated areas, if the sparsely populated areas are relatively insulated from interaction with the densely populated communities. In other words, the gravity model implies that linguistic features might not always spread in a geographically contiguous way, as one might first expect. The present paper seeks to evaluate the validity of the gravity model's assumption about the relation of interaction rates to linguistic diffusion rates, which we will refer to as the *interaction hypothesis*.

As for the actual mechanism by which linguistic features spread, the model adopts the functional-teleological view of chain shifts[4] and the exemplar theory (see Appendix A). Briefly summarized, it is assumed that individuals have a cache of short memories of vowels called *exemplars* which they draw on both to recognize an incoming acoustic signal as a given vowel and to map a mental vowel category onto an articulatory gesture that produces an outgoing acoustic signal. The cumulative interaction between individuals slowly modifies the cache of exemplars, leading to long-term changes in the individual's perception and production of vowels.

Because both the interaction hypothesis and the exemplar theory are essentially concerned with emergent phenomena, agent-based modeling is a natural way to investigate language contact within the current theoretical parameters.

The present model contains three communities of speakers, a densely populated `city`, a medium-density `town`, and a sparsely populated `village`. The inhabitants of each community are turtles referred to as `cityfolk`, `townsfolk`, `villagefolk`, respectively. The `cityfolk` and `townsfolk` travel along a "highway" directly connecting the two communities and which passes by the `village`; the `villagefolk` travel to either the `city` or the `town` depending on a fixed probability parameter `city-probability`. Every *time period* of 500 ticks, the turtles[5]

[4]This view assumes chain shifts occur to maintain clarity, akin to musicians adjusting pitch to avoid dissonance.

[5]In NetLogo, agents are called "turtles".

have a chance to begin traveling or return to their home community.

Each turtle stores a cache of vowels, which are modeled as single numbers. Although real vowels are characterized along multiple acoustic dimensions, this simplification follows the choice made by Stanford and Kenny, hence ensuring that our results are comparable to theirs. Furthermore, because vowels in a chain shift generally progress along a one-dimensional "track" in vowel space, one-dimensional vowels should still capture most of the qualitative phenomena.

When two turtles are on the same patch, they each speak three vowels to each other, and judge whether the vowels they heard are similar to their own cache of exemplars for that vowel category. If so, they add the new vowel token to their cache, which changes how they pronounce the vowel in the future. Through the continuous process of travel and speaking to other turtles, the collective vowel qualities for each community changes over time, modeling the progression of the chain shift.

3. Analysis and Results. We analyze the model outlined in § 2 in three parts. In § 3.1 we test the gravity model's prediction that a vowel shift can spread from one community to another while skipping sparsely populated intervening communities. In § 3.2 we test the gravity model's prediction that the rate of spread from one community to another increases with increasing population density of each community and decreasing distance between the communities. Finally, in § 3.3 we perform sensitivity analysis on the model to determine how robust its results are to different sets of parameter values.

Unless otherwise stated, we run each simulation in NetLogo's BehaviorSpace over various combinations of parameter values for 10,000 steps. This allows enough time for interesting phenomena to emerge without being unnecessarily computationally expensive.

3.1. Vowel category trajectories over time. To validate the gravity model's prediction that a change *can* spread from one community to another skipping over a third geographically intermediate community, we plot the three vowel category means of each community over

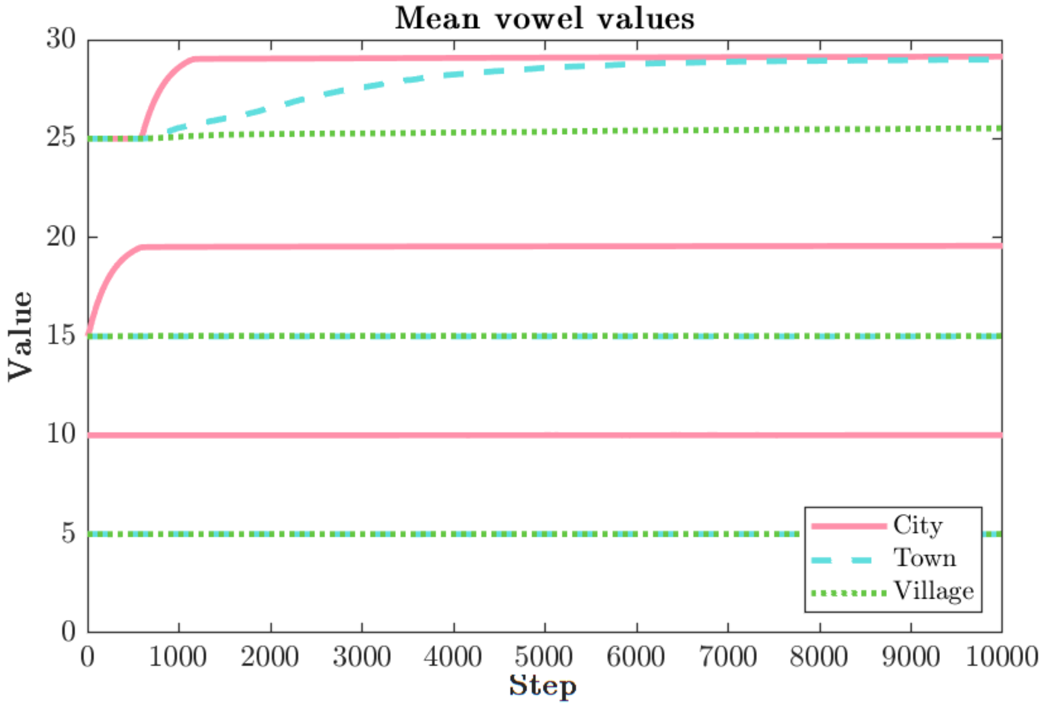

Figure 1. *A plot of all three communities vowel category means over time, averaged over 8 runs. Pink lines indicate the* `city`*'s vowel means, cyan lines the* `town`*'s, and green lines the* `village`*'s. The lowest line in each case is vowel A, the second highest vowel B, and the third highest vowel C. The village and town's vowels A and B differ so little no cyan line can be seen. For these runs,* `village-population` = 10, `village-radius` = 60, `village-distance` = 0, `city-probability` = 50.

time. If the prediction is true, then for at least one set of parameters, the `townsfolk`'s vowels should shift sooner than the `villagefolk`'s vowels. Moreover, we expect this prediction to be true when the population density of the `village` is much lower than that of the `town`, since the gravity model predicts this will lead the `village` to have overall fewer interactions with the `city` and hence a slower pace of change despite its closer distance.

Figure 1 shows a plot of the mean vowel values of the three communities over time, averaged over 8 runs. In these runs, `village-population` = 10, `village-radius` = 60, `village-distance` = 0, and `city-probability` = 50—the `village` is very sparsely populated. We first observe that the expected chain shift does occur in the `city`. Vowel A, being 5 units away from vowel B, is too close and causes vowel B to rise in value. The plot indicates that vowel B attains an equilibrium value just shy of 20 around 800 steps, and the model output verifies that vowel B changes by less than 0.0001 each step after step 640, attaining a value of 19.6 by the end of the simulation. This is too close to the `city`'s vowel C, which rises in turn to around 29 at about 1100 steps. The model output shows that the `city`'s vowel C changes are on the order of hundredths after about 1200 steps, eventually attaining a value of 21.9 by the end of the simulation.

Note in this case that the gravity model's prediction *is* borne out in this extreme case: we observe that the `village`'s green vowel C line rises more slowly than the `town`'s, indicating that the `village` adopts the change more slowly, and incompletely by the end of the 10,000 steps. The shift is likely slowed down by the village's sparser population and the fact that half the `villagefolk` travelers travel to the town instead of the city, as this is what distinguishes the `village` from the `town`.

In Figure 1, the `townsfolk` and `villagefolk`'s A and B vowels don't seem to change at all. This is likely due to the fact that a larger proportion of turtles are away from their home communities shortly after the start of the third time period (1000 ticks) compared to shortly after the start of the simulation. Since the `city`'s vowel C does not rise completely until well after the second time period has started, its values are closer to the `town`'s *after* the third influx of travelers, so that both communities' populations are more intermixed. At this point, there are both more interactions between `cityfolk` and `townsfolk` and less of a difference between their vowel systems, so conditions are much more favorable for the `townsfolk`'s vowel category C to converge to the `cityfolk`'s than they were for the `townsfolk`'s vowel category B to converge to the `cityfolk`'s. This shows that the shift's spread from one community to another is critically affected by the shift's progression in the community that drives the shift and the amount of mixing between the two community's populations. If the populations had started the model fairly mixed, it might be the case that all three vowel categories for the `town` would have shifted. A similar argument goes for the `village`, except the shift is slowed down even more by its sparser population and the fact that half the `villagefolk` travelers travel to the town instead of the city.

To summarize, the gravity model's prediction that a linguistic feature can spread from one community to another skipping an intervening community is verified at least when the village is quite sparse. Moreover, the shift's spread from one community to another is mediated by the shift's progression and the amount of mixture between the two communities throughout time.

3.2. Population density and distance. To validate the gravity model's prediction that a the rate of spread of a shift (or feature) from one community to another increases with increasing population density of the two communities and decreasing distance between them, we run the model and record the difference between the mean vowel C value for the `city` and the mean vowel C values for the `town` and the `village` while varying the parameters `village-population`, `village-radius`, and `village-distance`. Note that the village's population density should scale up with `village-population` and down with `village-radius`. Therefore if the gravity model's predictions are borne out, we should see the vowel differences decrease with increasing `village-population` and decreasing `village-radius` and `village-distance`.

To test the validity of this prediction, we first ran the model varying `village-population` from 10 to 80 in steps of 10 and `village-distance` from 0 to 50 in steps of 10, keeping `village-radius` fixed at 60 and `city-probability` fixed at 100. For each combination of parameters we ran the model 10 times for a total of $6 \times 8 \times 10 = 480$ runs. Figure 2 shows the resultant boxplots, with the `city`-`village` differences plotted in green and the `city`-`town` differences plotted in cyan. The medians and IQRs of mean vowel C differences are tabulated in Table 1.

The plots show a clear between-group difference, with the `village` medians being much higher than the `town` medians and the `village` interquartile ranges being generally larger

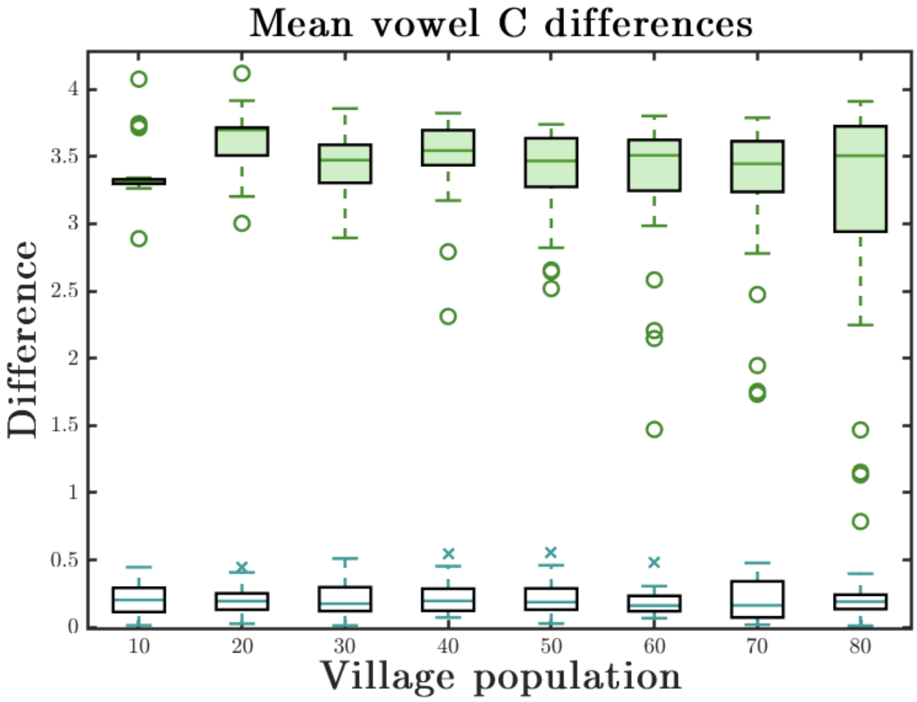

Figure 2. *Boxplots showing mean `city`-to-`village` (green, shaded, outliers marked with ∘) and `city`-to-`town` (cyan, unshaded, outliers ×) vowel C differences versus `village-population`. Values for `village-distance` are aggregated from 0–50. The value of `city-probability` is 100.*

		Mean vowel C differences							
Population		10	20	30	40	50	60	70	80
Village	Median	3.31	3.70	3.47	3.54	3.47	3.51	3.45	3.51
	IQR	0.03	0.21	0.28	0.26	0.36	0.38	0.38	0.78
Town	Median	0.20	0.19	0.17	0.20	0.19	0.16	0.16	0.19
	IQR	0.18	0.12	0.18	0.16	0.16	0.11	0.27	0.11

Table 1

Table of mean vowel C difference medians and IQRs displayed in Figure 2.

than the `town` interquartile ranges. Recalling that a smaller vowel difference corresponds to a more progressed shift, we observe that the `town` is more progressed in the shift than the `village` across `village-population` values. Recalling that `town-population` is set to 75 and that `town-radius` is set to 30 while `village-radius` is set to 60, this is in keeping with the gravity model's predictions—despite the `village`'s closer proximity to the `city` and the fact that all `villagefolk` travel to the `city`, its lower population density apparently causes it to pick up on the shift slower.

On the other hand, there is no significant within-group difference for either the `village` or the `town` difference medians across `village-population` values. This is expected for the `town` since any effects of `village-population` on the `townsfolk` would be indirect, but unexpected for the `villagefolk`—the gravity model would predict that we see smaller vowel differences with increasing `village-population`, as this increases the `village`'s population density. We do however observe increased interquartile means in the direction of lower differences, suggesting that there is a weak effect of increasing `village-population` on the progress of the change.

We next repeated the 480 runs with the same set of parameter values, but with `city-probability` set to 50. The resulting boxplots are displayed in Figure 3; the medians and IQRs are tabulated in Table 2. For these runs, we observe the same between-group difference, where the `town` is much farther progressed than the `village` across all `village-population` parameter values. However, we now observe a within-group trend for the `villagefolk`: the median vowel differences generally decrease with increasing `village-population`, as the gravity model predicts. At a `village-population` of 80, the median difference is just below 1.5, much lower than the corresponding median difference of about 3.5 in the runs with `city-probability` set to 100, indicating that the shift generally progresses much farther at this lower `city-probability` value. This is highly unexpected, since the only difference between these runs and the previous runs is that `city-probability` is decreased to 50, so that the `villagefolk` visit the `city` *less* and have *fewer* interactions with the `cityfolk`.

Finally, note in Figure 3 the small dip at very low values of `village-population`. One possible explanation is that at very low populations, a single `villageperson` who has assimilated to the `city`'s vowel trends will bring the `village` vowel C mean much closer to the `city` vowel C mean. If this is the case, shifts progress faster in smaller communities in a purely numerical sense—there are fewer users to acquire the shift in the first place, and so using a

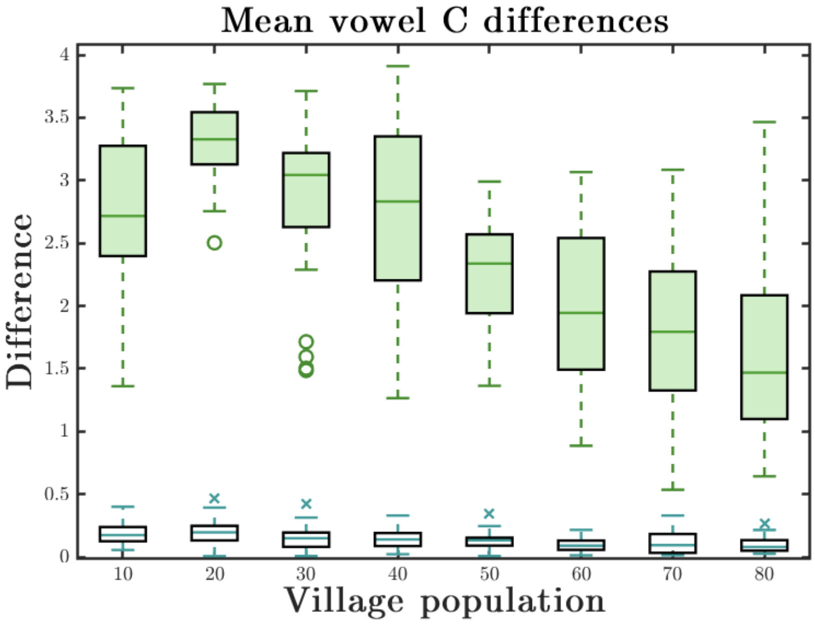

Figure 3. *Boxplots showing mean* `city`*-to-*`village` *(green, shaded, outliers marked with ○) and* `city`*-to-*`town` *(cyan, unshaded, outliers marked with ×) vowel C differences versus* `village-population`*. Values for* `village-distance` *are aggregated from 0–50. The value of* `city-probability` *is 50. Medians are labeled on the boxes and IQRs are labeled to the right of each box.*

		Mean vowel C differences							
Population		10	20	30	40	50	60	70	80
Village	Median	2.72	3.33	3.04	2.83	2.34	1.94	1.79	1.47
	IQR	0.88	0.42	0.59	1.15	0.63	1.05	0.95	0.99
Town	Median	0.17	0.20	0.15	0.14	0.13	0.09	0.09	0.08
	IQR	0.11	0.12	0.11	0.10	0.06	0.07	0.15	0.08

Table 2

Table of mean vowel C difference medians and IQRs displayed in Figure 3.

summary statistic like the mean means the shift will appear to progress much faster. These results suggest there may be a population density threshold under which changes diffuse faster through *sparser* rather than denser communities. We leave follow-up on this idea to future research.

Similar patterns appear when the parameter varied is `village-radius`. Under the gravity model, we would predict that a larger radius, representing a lower population density, is correlated with a slower progression of the shift.

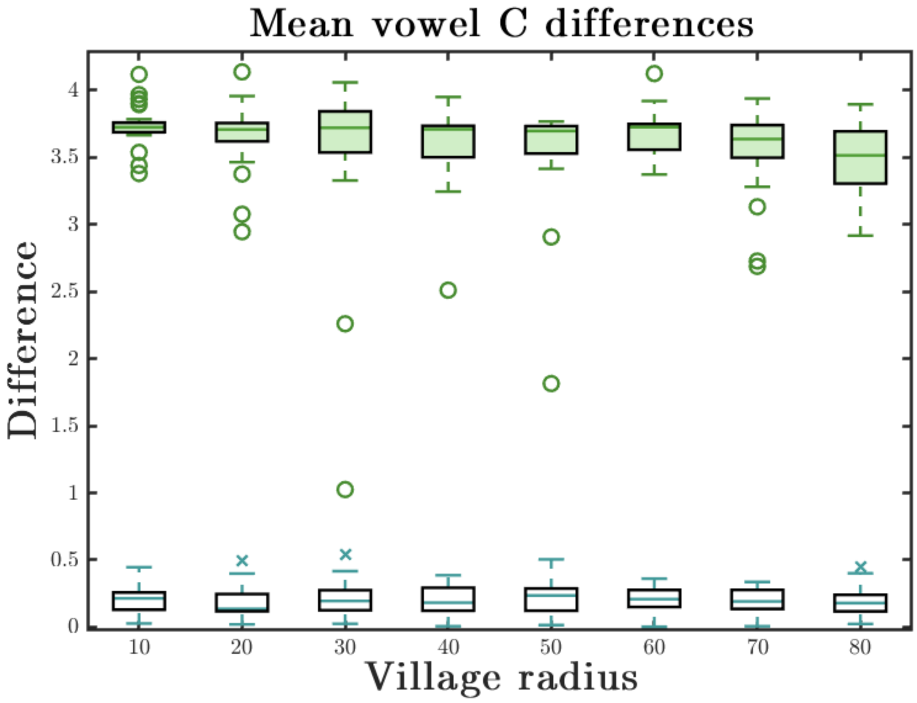

Figure 4. *Boxplots showing mean `city`-to-`village` (green, shaded, outliers marked with ◦) and `city`-to-`town` (cyan, unshaded, outliers marked with ×) vowel C differences versus `village-radius`. Values for `village-distance` are aggregated from 0–50. The value of `city-probability` is 100.*

		Mean vowel C differences							
Radius		10	20	30	40	50	60	70	80
Village	Median	3.72	3.71	3.72	3.71	3.70	3.73	3.64	3.51
	IQR	0.07	0.14	0.31	0.23	0.20	0.19	0.24	0.39
Town	Median	0.21	0.14	0.20	0.18	0.24	0.21	0.19	0.18
	IQR	0.13	0.13	0.15	0.17	0.17	0.13	0.14	0.12

Table 3

Table of mean vowel C difference medians and IQRs displayed in Figure 4.

To investigate the impact of `village-radius` on the progression of the shift, we varied `village-radius` from 10 to 80 in steps of 10 and `village-distance` from 0 to 50 in steps of 10, keeping `village-population` fixed at 20 and `city-probability` fixed at 100. Each parameter combination was simulated for 10 runs, again totaling $6 \times 8 \times 10 = 480$ runs. Boxplots with the `city`-`village` differences plotted in green and the `city`-`town` differences plotted in cyan appear in Figure 4; the relevant medians and IQRs are tabulated in Table 3.

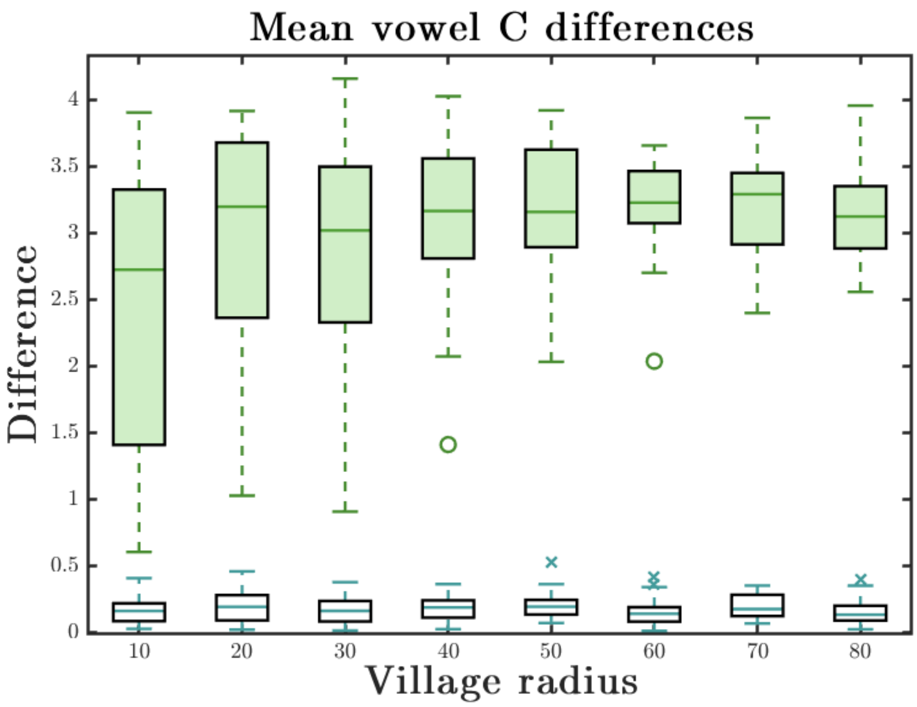

Figure 5. *Boxplots showing mean `city`-to-`village` (green, shaded, outliers marked with ∘) and `city`-to-`town` (cyan, unshaded, outliers marked with ×) vowel C differences versus `village-radius`. Values for `village-distance` are aggregated from 0–50. The value of `city-probability` is 50.*

		Mean vowel C differences							
Radius		10	20	30	40	50	60	70	80
Village	Median	2.72	3.20	3.02	3.16	3.16	3.23	3.29	3.12
	IQR	1.92	1.32	1.17	0.75	0.73	0.39	0.54	0.47
Town	Median	0.16	0.19	0.16	0.18	0.19	0.14	0.17	0.13
	IQR	0.13	0.19	0.15	0.13	0.11	0.11	0.16	0.11

Table 4

Table of mean vowel C difference medians and IQRs displayed in Figure 5.

We observe the same pattern as we did with Figure 2, where there is a clear between-group difference but no within-group difference for the `villagefolk`. Moreover, we do not observe the trend of increasing IQRs with decreasing `village-radius` as we had with increasing `village-population`. If there is any effect of population density in this case, it is exceedingly weak.

Next we repeated the runs with the same parameters but with `city-probability` set to

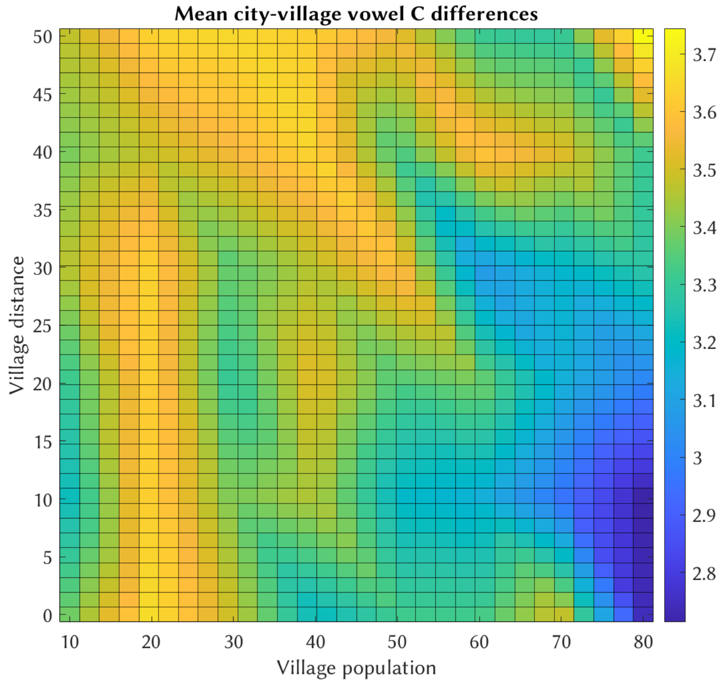

Figure 6. *Heatmap showing mean `city`-to-`village` vowel C differences against `village-population` and `village-distance`. The value of `city-probability` is 100.*

50. The results are plotted in Figure 5, and the associated medians and IQRs are tabulated in Table 4. We observe a slight downward trend in the `village` median vowel differences with decreasing `village-radius` along with a dramatic increase in IQR and a shift in the direction of 0. This is the expected pattern predicted by the gravity model, but it again occurs at a lower `city-probability` value where the `villagefolk` interact *less* with the `cityfolk`.

Next, we investigate the impact of the parameter `village-distance` on the progression of the shift. The same data used to plot Figure 2 was used to make the heatmap in Figure 6, except the `city`-`town` differences are not plotted and the outcomes for each value of `village-difference` is plotted separately.

Note first that there is not as clean and simple a pattern as the gravity model might lead one to expect. There are "hills" and "valleys" in the mean vowel C differences, so that the dependence of shift progression on both `village-population` and `village-distance` is not monotonic. Despite this, general trends can be noted. For instance, at high values

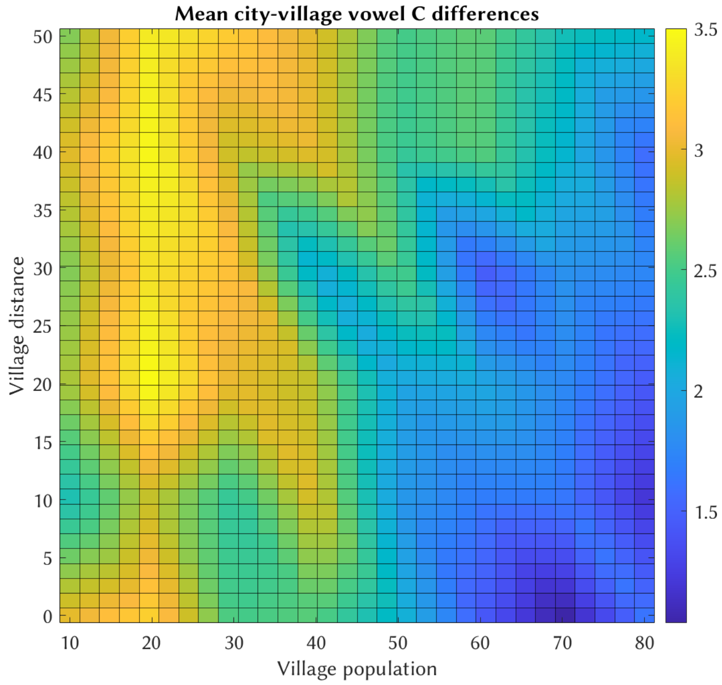

Figure 7. *Heatmap showing mean `city`-to-`village` vowel C differences against `village-population` and `village-distance`. The value of `city-probability` is 50.*

of `village-distance` (over 33), most of the heatmap is yellow, indicating relatively high differences of around 3.6 or above. That is, when the `village` is far from both of the other communities, the shift does not progress very far. This is the result we expect under the gravity model. Similarly, at high values of `village-population` (over 72), we observe that most of the heatmap colors at a relatively low value of `village-distance` (under about 26) are blue, indicating smaller difference values of 2.8–3.2 and a slightly faster shift progression. Again, this is expected under the gravity model. However, notice that even the minimum difference value is 2.8, so that when `city-probability` = 100 the change never progresses very far. Hence much of the nonmonotonicity may be able to be explained by noise from stochasticity in the model.

Figure 7 is the analog of Figure 6, using the same data as previously plotted in Figure 3. In this case the trends previously noted are much more apparent. Most of the heatmap boxes at a `village-population` value higher than 50 are green or blue, representing a difference of 1–

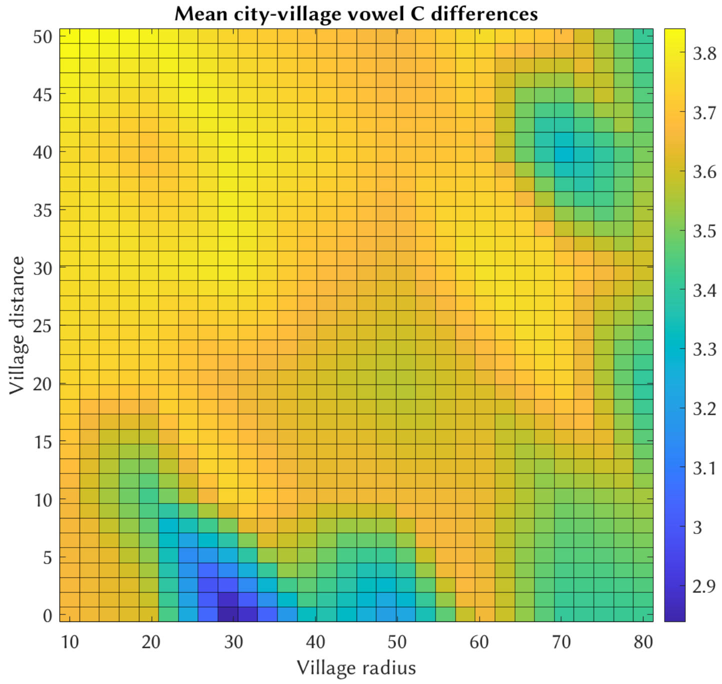

Figure 8. *Heatmap showing mean* `city`*-to-*`village` *vowel C differences against* `village-radius` *and* `village-distance`*. The value of* `city-probability` *is 100.*

2.5, and most of the heatmap boxes under this value are orange or yellow (a difference of 3–3.5) given a `village-distance` of 15 or above. These results match the gravity model's predictions much better, and we see a lower minimum difference, showing that when `city-probability` is 50, the shift progresses farther for the `village`.

Similar results obtain when `village-radius` is varied instead of `village-population`. Figure 8 shows the `city`-`village` data of Figure 4 plotted as a heatmap, with `city-probability` at 100. Note the extraordinarily small range of differences (2.85–3.85) and the unexpected observation that smaller differences (a more progressed shift) result for larger values of `village-radius` (a sparser `village`). Again, it is likely that much of the variation here can be explained by inherent stochasticity in the model.

Figure 9 shows the `city`-`village` data of Figure 5 as a heatmap, with `city-probability` at 50. Here a clearer pattern emerges; there appears to be a linear boundary that runs from a `village-radius` value of 10 and a `village-distance` value of 40 down to a `village-radius`

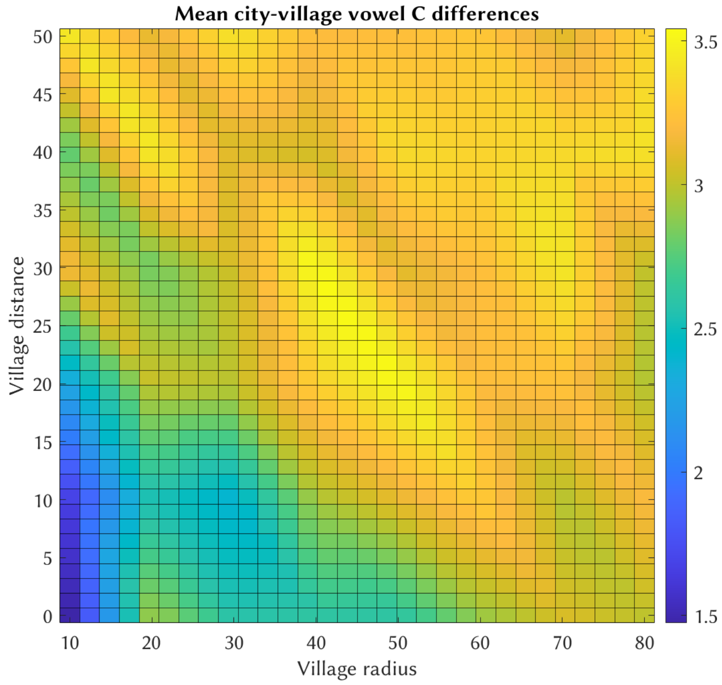

Figure 9. *Heatmap showing mean* `city`*-to-*`village` *vowel C differences against* `village-radius` *and* `village-distance`*. The value of* `city-probability` *is 50.*

value of 55 and a `village-distance` value of 0. To the left of this boundary, the heatmap is generally green or blue (differences of 1.5–2.7, a well-progressed shift). To the right of this boundary, the heatmap is generally orange or yellow (differences of 3.3–3.5, not much shift progression). This boundary suggests that the shift progresses farther for smaller values of `village-radius` and smaller values of `village-distance`, which accord with the gravity model's predictions that a denser population and a closer distance should result in a more progressed shift.

To summarize, the previous results show that the gravity model's predictions are not borne out when `city-probability` is 100, but that they are when `city-probability` is 50. This is despite the fact that `villagefolk` interact with the `cityfolk` much more at higher `city-probability` values. We can explain this counterintuitive result by considering the interactions between `villagefolk` and `townsfolk`.

To investigate the effect of `city-probability` on the progression of the shift in the

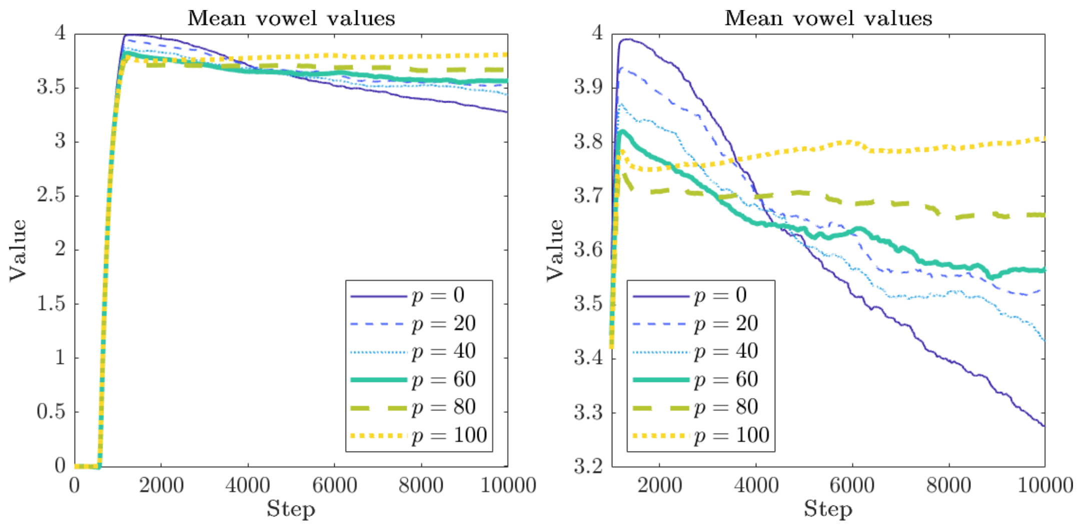

Figure 10. *Plots of city-village vowel C differences over time for different values of* **city-probability**, *represented as p in the legend. Left: full results of simulation. Right: detail of timesteps 1000–10000.*

village, we ran the simulation with village-distance fixed at 100, village-population fixed at 20, and village-radius fixed at 60, varying the value of city-probability from 0 to 100 in steps of 10. Each parameter value was carried out for 10 runs, and the resulting city-village vowel C differences were averaged and plotted over the step number. The result is displayed in Figure 10.

Analyzing this plot, we observe that until around 750 ticks or so there is no difference between the city and village's vowel C. Comparing this plot with Figure 1, we see that this is likely because the cityfolk's vowel B has not yet risen high enough to push the city's vowel C up. At this stage, both the city and the village vowel C categories are at their initial value of 25. At this stage, the value of city-probability has not resulted in any different behavior, as can be seen from the variously-colored lines overlapping nearly perfectly. As the city vowel C rises, the difference between the two steadily grows, until around about 1000 ticks. In a narrow interval around this point, all the trajectories hit their peaks. The higher city-probability trajectories hit lower difference peaks, then slowly rise in value. On the other hand, the lower city-probability trajectories hit higher difference

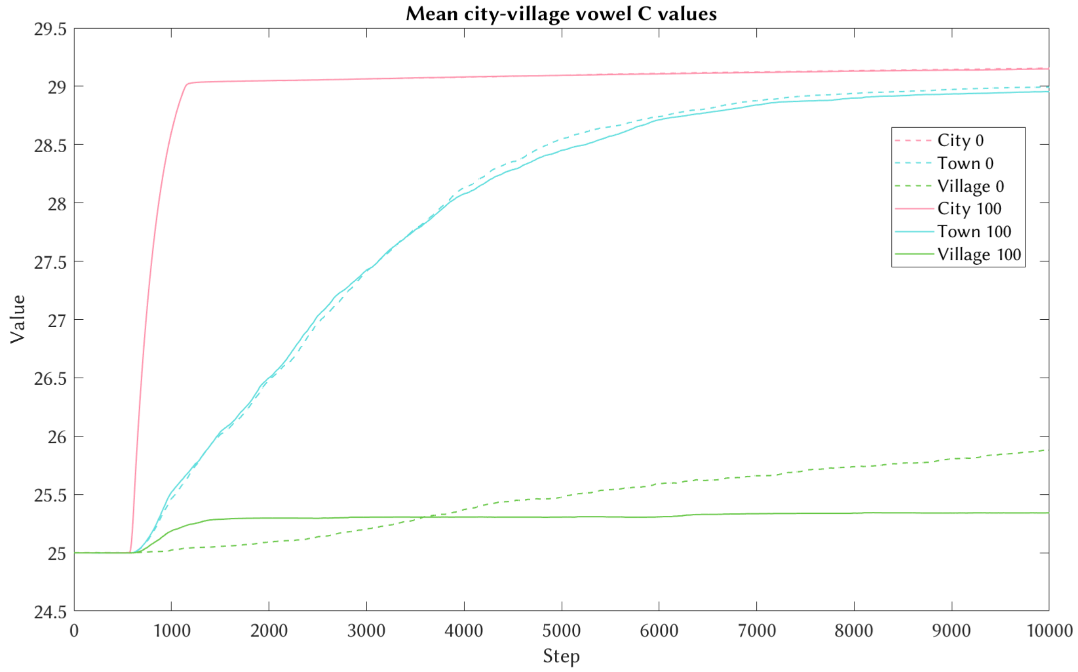

Figure 11. *Plots of `city`, `town`, and `village` vowel C values over time for `city-probability` values of 0 (dashed line) and 100 (solid line).*

peaks, but then steadily decrease until by 10,000 ticks they reach lower differences than did the high-probability trajectories.

We can get a clearer view of this by plotting the three communities' vowel C values separately for `city-probability` values of 0 and 100. Figure 11 shows just such a plot. In the high-probability case, the `village` trajectory flattens out after a while, producing the flat trajectories seen in Figure 10. In the low-probability case, the `village` trajectory starts rising later but continues to increase throughout the 10,000 steps, producing the dip seen in Figure 10. The `city` and `town` trajectories do not vary significantly in either scenario.

Recall that if a traveling `villageperson` does not travel to the `city`, they travel to the `town`. Figure 11 shows clearly that the `town` acquires the shift much more gradually than the `city`, and hence stays closer to the `villagefolk`'s category C. If in this case the `townsfolk` with lower category C values than the mean speak to `villagefolk`, it is possible that their category C values are accepted by the `villagefolk`. On the other hand, the `cityfolk`'s vowel category C is so different `villagefolk` do not accept it as an instance of their own category and do not converge. In other words, the model predicts that increased interaction rates do not always lead to higher rates of diffusion if the differences between the two interacting populations is too great, since members of one community may not map the acoustic signals produced by another onto their own mental categories perfectly. Although it is not observed here, we may also expect that in realistic situations, saturation effects may occur, such that past a certain threshold increased interaction rates no longer lead to significant increases in diffusion rate.

In summary, the above analysis shows that the gravity model's predictions hold true only in a subset of situations, and that this subset is determined by interaction patterns and the `tolerance` thresholds at which turtles will accept novel vowel tokens as instances of their own categories. In particular, the gravity model's predictions will hold true in situations where two interacting communities have categories within each other's tolerance thresholds. Because `tolerance` and other parameters were arbitrarily set in the preceding analysis, in the next section we analyze how robust the model's output is to changes in parameters.

3.3. Sensitivity analysis. In the previous section, we showed that `tolerance` was an important parameter in determining how and to what extent the shift spreads across the three communities. There are a number of other parameters in the model that have been arbitrarily set, generally following the values of [7] or scaling them down where possible to reduce computational complexity. However, it is not clear that these parameter values are the "right" ones that correspond to real-world situations. In some cases, such as that of `memory-size`, it is clear that they cannot correspond directly to real-world situations, potentially putting the usefulness of the model results in interpreting real-world phenomena in question.

To address this problem, this section uses [6]'s Latin hypercube sampling/partial rank correlation coefficient (LHS/PRCC) strategy for conducting sensitivity analysis efficiently on many parameters at a time. In brief, the PRCC is a robust measure of variable monotonicity

Parameter	Minimum	Mean	Max
`memory-size`	10	30	50*
`tolerance`	0.01	1*	4
`drift-rate`	0.01	0.1*	1
`max-stay`	1	4*	7
`percent-travelers`	0.01	0.05*	0.1
`initial-dev`	0.01	0.1*	1

Table 5

Parameters used in the LHS/PRCC analysis and the inputs to the procedure. Starred values indicate the values used in the preceding analyses.

after removing linear relationships in the data. By using LHS to efficiently sample the parameter space, it is possible to test which parameters affect the model output (progression of the shift in the `village`, operationalized by mean `city`-`village` vowel C difference) significantly. If it is shown that the specific values of the parameters do not significantly affect the model output, then we can have some confidence that our results can be safely transferred to real-world situations.

Table 5 shows the parameters used and inputs to the LHS/PRCC procedure. Figure 12 shows the outputs of the LHS/PRCC procedure. The clear outlier is `tolerance`, which has by far the highest PRCC magnitude and the lowest p-value. At a significance level of $\alpha = 0.1$, only `tolerance` significantly affects the model output. This suggests that the results obtained by the analyses above are generally robust to parameter changes, but that the exact exemplar tolerance thresholds may need to investigated further.

4. Conclusions. This paper has developed an agent-based model to test the validity of the gravity theory of linguistic diffusion. The model uses an exemplar model and the functional-teleological explanation of chain shifts to build up a picture of how individual agents change their vowel systems, and studies the emergent patterns across language communities. It finds that the predictions of the gravity model hold true in a subset of cases determined at least in part by patterns of interaction and the particularities of how vowels are produced and perceived. To be specific, the gravity model's predictions will hold true in situations where

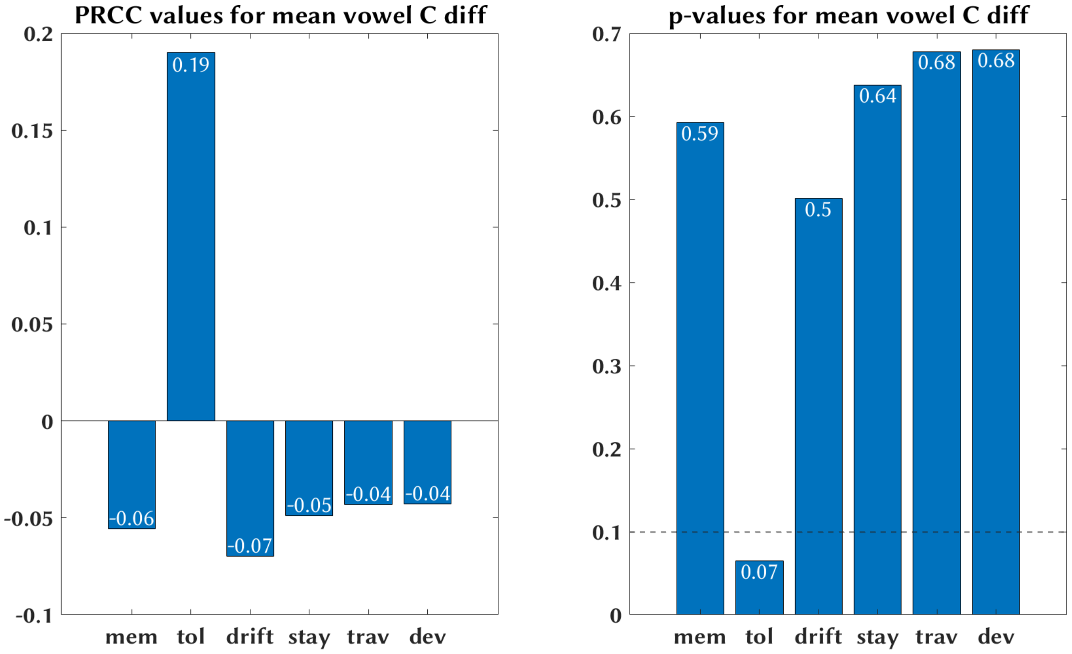

Figure 12. *PRCC and p-values for the mean* `city`-`village` *vowel C difference. The parameters changed and their abbreviations are* `memory-size` *(mem),* `tolerance` *(tol),* `drift-rate` *(drift),* `max-stay` *(stay),* `percent-travelers` *(trav),* `initial-dev` *(deviation). The* $\alpha = 0.1$ *significance level is represented as a dotted horizontal line on the p-value plot.*

two interacting communities have categories within each other's tolerance thresholds.

The present model is limited in both its complexity and simplicity. Theoretically, it would be desirable to arrive at a more economical model with fewer submodels and fewer parameters. However, the present model also makes a number of simplifying assumptions that are not always justified, such as those about the movement of language users and the ways in which they produce and perceive vowels. However, the sensitivity analysis in § 3.3 showed that the model results were robust to changes in parameter values, suggesting that the results may be generalizable to real chain shifts with many different particularities.

The present study has modeled the case of just one sparse community between two larger, denser communities. It is more realistic to consider several small communities, each not necessarily sparse when considered alone, but collectively sparse over the area between the two larger communities—that is, a smattering of `villages` between the `city` and the `town` instead of just one. This would correspond more closely to the commonplace situation of many small towns off the highway between two larger cities. The present model has simplified the vowel space into a one-dimensional line. Alternatively, it would also be possible to model the vowels as points in a two- or multidimensional space. In this case, the different sensitivities of speakers to different acoustic cues (F1, F2) might lead to different patterns in the spread of the shift. The present model is concerned with diffusion through space, but it could also be adapted to a social network rather than a spatial grid. In this case, it would be interesting to see how shifts spread from locally well-connected subnetworks to the larger network.

REFERENCES

[1] N. J. Enfield, *Areal linguistics and mainland Southeast Asia*, Annu. Rev. Anthropol., 34 (2005), pp. 181–206.

[2] V. Grimm, *The ODD protocol: An update with guidance to support wider and more consistent use*, Ecological Modelling, 428 (2020), 109105, p. 109105, https://doi.org/10.1016/j.ecolmodel.2020.109105.

[3] V. Grimm and S. F. Railsback, *Designing, formulating, and communicating agent-based models*, in Agent-based models of geographical systems, Springer, 2011, pp. 361–377.

[4] M. Haspelmath, *How young is standard average European*, Language sciences, 20 (1998), pp. 271–287.

[5] W. Labov, *The atlas of North American English: Phonetics, phonology, and sound change*, Walter de Gruyter, Berlin, 2008.

[6] S. Marino, I. B. Hogue, C. J. Ray, and D. E. Kirschner, *A methodology for performing global uncertainty and sensitivity analysis in systems biology*, Journal of theoretical biology, 254 (2008), pp. 178–196.

[7] J. N. Stanford and L. A. Kenny, *Revisiting transmission and diffusion: An agent-based model of vowel chain shifts across large communities*, Language variation and change, 25 (2013), pp. 119–153.

[8] S. G. Thomason and T. Kaufman, *Language contact*, vol. 22, Edinburgh University Press Edinburgh, 2001.

[9] W. Wolfram and N. Schilling-Estes, *Dialectology and linguistic diffusion*, The handbook of historical linguistics, (2017), pp. 713–735.