Computational Linguistics in the Netherlands Journal 11 (2021) 5-26 Submitted 09/2021; Published 12/2021

Keeping up with the Neighbours - An Agent-Based Simulation of the Divergence of the Standard Dutch Pronunciations in the Netherlands and Belgium

Anthe Sevenants* ANTHE.SEVENANTS@STUDENT.KULEUVEN.BE
Dirk Speelman* DIRK.SPEELMAN@KULEUVEN.BE

**KU Leuven, Leuven, Belgium*

Abstract

While the Netherlandic standard Dutch pronunciation norm around 1930 was still very much like the Belgian norm, it shifted considerably in the course of the 20th century (Van de Velde 1996, Van de Velde et al. 2010). In Belgium, no such evolution occurred, which caused the pronunciation of both language varieties to diverge. As of yet, there is no conclusive evidence as to why this divergence has happened. Because there is not enough data to investigate the divergence empirically, it is examined using an agent-based simulation model in Python. Though we cannot 'prove' that the mechanisms described in the theories from the literature actually happened in reality, we *can* test their plausibility by checking whether the effects described in the theories also appear in our model which attempts to mimic real-world circumstances. Four research questions based on theories found in the literature are tested: 1. Can a reduced contact between speakers from the Netherlands and Belgium result in a divergence between the standard pronunciations of both countries in the model? 2. Can an increased pace of language change in Dutch speakers cause a divergence between the standard pronunciations of the Netherlands and Belgium in the model? 3. Can we relate increased ethnocentrism in Belgian speakers to less adoption of Netherlandic innovations or even divergence in the model? 4. Can an increased media influence amplify the existing tendencies for language change (acceleration or inhibition) in Belgium in the model? The results show that a lack of contact between both countries can indeed lead to divergence in the model, but only if abroad travel is at least 5000 times less likely than domestic travel. The pace of language change in the Netherlands does not have a sizeable impact on convergence or divergence tendencies in Belgium in the model. High values for ethnocentrism in Belgian agents are able to lead to divergence in the model, as long as these high values are shared by the entire population. If ethnocentrism decreases along with how close agents live to the border, it has little effect. Media receptiveness in agents always kickstarts convergence in the model and it accelerates this convergence as well. Since media influence is implemented as a powerful force in the simulation, this result must be interpreted from the viewpoint of media having a sizeable impact on language change.

1. Introduction

Computer simulations are a booming research field in linguistics. They are increasingly used to simulate the development, spread and exchange of language in areas where researchers lack experimental data. In the field of sociolinguistics, this lack of data emanates from a paradoxical *abundance* of data. While the necessary data is — in principle — out there, it is organisationally infeasible to conduct the necessary experimental research to collect this data, not to mention that it would be prohibitively expensive. It is often impossible for a researcher or even a group of researchers to follow a large body of participants extensively and record their language use (Kretzschmar et al. 2014, Beuls and van Trijp 2016).

An additional problem for the study of language change as a whole is the fact that language change is not a centrally organised process. Kretzschmar et al. (2014) warn that theoretical concepts such as Neogrammarian sound laws or lexical diffusion present language change as a mechanical process, which downplays the important role of individual speaker innovations that drive these changes. They prefer to look at language as a 'Complex Adaptive System' (CAS), a term also echoed in other simulationist research (Kirby 1999, Gong and Shuai 2013, Smith 2014). Smith (2014, p. 284) explains the term as follows:

> A complex adaptive system is complex because the behavior of the system as a whole cannot be projected trivially from the microbehavior of its individual components, and adaptive because this microbehavior changes and develops in response to historical experience and external pressures from the wider environment.

Language is a Complex Adaptive System because different speakers adapt their language use independently from each other while at the same time copying others. Eventually, this leads to a change in the 'language system' as a whole, though the changes originated from the individual speakers who — either consciously or subconsciously — reshaped their linguistic behaviour (Kretzschmar et al. 2014).

This complex behaviour of a CAS gives rise to a methodological problem for typical empirical research. For sociolinguists, it would be strictly infeasible to assess the influences of individual speaker behaviour on a language system, especially if this behaviour changes along with the system. As a result, computer simulation becomes 'the only way that we can model diffusion as the adaptive aspect of complex systems in speech and culture' (Kretzschmar et al. 2014). Simulations allow researchers to have thousands of virtual speakers or speaker communities — also called 'agents' — talk to each other on a very local level. Researchers can then check what effect the local behaviour has on the system as a whole. In other words, the CAS principle can be virtualised in order to supply researchers with scientifically backed data to test their principles in a controlled environment. At the same time, we can minutely track how, with whom, and when agents speak. Of course, computer simulations are also not limited by time and space; one can perfectly test a hypothesis dating back thousands of years (Gong and Shuai 2013).

It is important to note that simulation research does not aim to find definitive answers to research questions; simulations are not omniscient beings. They are mere *exploratory* tools which aim to examine whether a certain process *could* be responsible for a certain phenomenon in the real world (Pijpops and Beuls 2015, Gong et al. 2014). The train of thought is that 'if [certain] phenomena emerge in [a] computer model, then we will be able to suggest that such underlying principles may operate in the real world as well' (Stanford and Kenny 2013). Still, we can never be completely sure whether a model aligns with reality: all simulations do is confirm or deny plausibility within the model architecture. If we assume that this model architecture is accurate, however, the results can be very informative. These results (what we *actually* want to find out from the model) stem from an inquiry into the minimal conditions needed to give rise to the behaviour at hand (Baronchelli 2014). Approaching the problem in this way implies that, in order to find the minimal conditions, we try to consider a very wide range of theoretically possible model assumptions. These need not necessarily come from external sources (they need not be 'indisputable givens'), since the entire purpose of a simulation study is to find the circumstances which allow for a specific phenomenon, but also which circumstances do not. This interpretation of what a model can be is radically different from the traditional way of looking at models. For example, to design a Hidden Markov Model, the first step would be to extract the appropriate probabilities from a dataset. A simulation model reverses this process: rather than inducing a model from the data, we try to consider a wide range of theoretically possible models to find out which range of models can give rise to our phenomenon under scrutiny. It is this range (and what conditions the range) that is of particular interest in simulation studies. The term 'plausibility', then, refers to a 'possibility', 'workability' or the success of a proof of concept. It does not relate to the phenomenon at hand in a direct sense (in contrast to the interpretation of this term in traditional empirical research).

Because of their power, simulation studies have found their way into many different subbranches of linguistics. There are simulation studies in the field of language evolution (e.g. Beuls and Steels 2013), cognitive linguistics (e.g. Kirby 1999), historical linguistics (e.g. Pijpops et al. 2015) and of course sociolinguistics (e.g. Pierrehumbert et al. 2014). Because the case study further in the article is on the intersection of historical linguistics and sociolinguistics, we will briefly discuss one study from each of these fields.

In the field of historical linguistics, Pijpops et al. (2015) examined in their simulation study how the weak verbal conjugation system in Germanic was able to become dominant. This is an especially interesting question considering that (1) the strong verbal conjugation system used to be much more regular than we now perceive strong verbal systems to be in present-day Germanic languages (2) the weak verbal conjugation system had to start from a position of zero use. By testing a range of theoretical assumptions (cfr. supra), the researchers were able to find that the assumption of general applicability of the weak verbal system was sufficient to model its rise to becoming the dominant system. The outcome of their simulation only makes a *suggestion* about the actual plausibility of this conclusion, but it is theoretically interesting to examine these verbal systems in a synthetic environment.

In the field of sociolinguistics, Pierrehumbert et al. (2014) built a model to investigate how an innovative language form can propagate through a social network. The model actually tests many assumptions, but only the assumption that closeness to an innovation is conducive to the adoption of that innovation was able to produce a diffusion in the model with realistic properties. This again goes to show how important the range of assumptions in the model really is, since not every assumption necessarily produces the desired outcome.

2. Case study: divergence of the standard Dutch pronunciations in the Netherlands and Belgium

2.1 Introduction

One problem in historical sociolinguistics which lends itself to the use of computer simulations is the divergence of the standard Dutch pronunciations in the Netherlands and Belgium. At the beginning of the 20th century, the standard language pronunciations in The Netherlands and Belgium were more or less alike (Van de Velde 1996). This was the direct result of a decision by Flemish speakers to import the Dutch language norm on the basis of status and prestige (van der Sijs and Willemyns 2009). Flanders desperately needed a unified language standard with status and prestige to pit against French, which was steadily becoming more and more important in Belgium as a whole. The northern Dutch norm, associated with the glory periods of the Kingdom of the Netherlands, was a perfect fit for this.

Shortly after this brief period of pronunciation parallelism, however, the standard pronunciation in the Netherlands started to shift around the 1930s (Van de Velde 1996). The most notable aspects of this shift were, among others, the diphthongisation of the [e] and [o] sounds towards [e^{i}] and [o^{i}] and the devoicing of syllable-initial fricatives.[1] In Belgium, the standard Dutch pronunciation did not change. Speakers from Flanders did not follow the northern innovations, which caused the Dutch and Belgian pronunciation standards to diverge. Ironically, this also means that the Belgian-Dutch pronunciation standard, which speakers from the Netherlands nowadays regard as 'typically Belgian', is in fact simply an older pronunciation stage of the standard language in the Netherlands, but it is no longer perceived as such (van Oostendorp 2016).

What interests us in this specific case is the motivation behind *why* speakers of Dutch in Belgium did not adapt their standard language norm to the new innovations from the Netherlands. As of yet, there is no evidence as to why the divergence between the Dutch standard pronunciations in the Netherlands and Belgium happened, and because of a lack of data (cfr. supra), there probably

1. For an elaborate discussion, see Van de Velde (1996) and Van de Velde et al. (2010).

never will be. A simulation study of the divergence is interesting since it allows us to test the ideas that researchers have had about this divergence over the past decades. By implementing these ideas as assumptions in a series of models, we can see how they perform in a synthetic computer simulation context. Again, it is impossible to prove through a simulation model that a certain theory was indeed responsible for the divergence — if we had the data to prove this, we would not use simulation models. Rather, we try to find out which operationalisations of the theories 'work' under which (synthetic) circumstances, but we can never prove that they were indeed responsible as well.

As was mentioned before, several theories have been posited in the literature in an attempt to explain the divergence. A first theory states that there was simply not enough contact between Belgium and the Netherlands for Belgian speakers to follow the linguistic innovations in the Netherlands (van den Toorn 1997, Deprez 1985). As a result, an indigenous standard pronunciation spontaneously arose in Belgium. A second theory concerns the pace at which language changed in the Netherlands. According to Van de Velde (2019), the standard pronunciation in the Netherlands evolved too fast for Belgian speakers to adopt these innovations. A third theory relates to the attitude of Belgian speakers towards speakers from the Netherlands. Van de Velde (1996) refers to van Istendael (2005) and Deprez (1985), who state that Belgian speakers simply did not want to sound like speakers from the north. The problem with this explanation is that attitudinal research towards Netherlandic Dutch, such as the research which Deprez refers to, only appeared in the late 1970s and the 1980s. This is a period in which Belgium, and most notably Flanders, started becoming culturally independent and less reliant on the Netherlands (Geeraerts 2017). It is therefore uncertain whether we can 'retroactively' apply language attitudes from the 1980s to the situation in the 1930s, but it remains an interesting consideration nonetheless. A final theory concerns the role of the public broadcasting corporations in both countries. According to Van de Velde (1996), the Dutch broadcasting corporation has historically been more open to linguistic innovations, while the Belgian radio and television have traditionally been very conservative. It is possible that the Dutch laissez-faire approach accelerated ongoing language innovations, while the Belgian conservative approach actively inhibited such innovations. In short, these theories leave us with four general research questions:

1. Can a reduced contact between speakers from the Netherlands and Belgium result in a divergence between the standard pronunciations of both countries in the model? If so, where could the tipping point between convergence and divergence lie?

2. Can an increased pace of language change in Dutch speakers cause a divergence between the standard pronunciations of the Netherlands and Belgium in the model? If so, where would the tipping point lie?

3. Can we relate increased ethnocentrism in Belgian speakers to less adoption of Netherlandic innovations or even divergence in the model? If so, where would be the tipping point?

4. Can an increased media influence amplify the existing tendencies for language change (acceleration or inhibition) in Belgium in the model? If so, could there be a tipping point?

2.2 Model description

2.2.1 Space

Our model attempts to simulate cross-country communication in nine cities across the Netherlands and Belgium, mixing central, less central, border and non-border cities such as Antwerp, Utrecht, Mechelen and Tilburg. The model is implemented in the Python programming language and was built on top of the MESA simulation library (Kazil et al. 2019), which provides a ready-made infrastructure to ease the design process of computer simulations.[2]

The model presented here is loosely based on Stanford and Kenny (2013), who modelled traffic between Chicago and Saint Louis as part of their simulation. Like in Stanford and Kenny (2013), the model world consists of a grid with cells which agents can occupy. Agents are not allowed to occupy the space between two cells, but one cell can be occupied by multiple agents at once.

Also on the spatial dimension, the model makes extensive use of influence spheres, as in Stanford and Kenny (2013), which more or less represent the influence radius of cities. In our model, an influence sphere is established for every city under scrutiny at a realistic distance from the other spheres. Real-world distances could not be used because they would complicate the model design too much. Nine cities are represented in the model: for the Netherlands, these cities are Rotterdam, Hilversum, Breda and Tilburg, while for Belgium, Antwerp, Turnhout, Mechelen, Leuven and Dendermonde are included. The cities were specifically chosen to include linguistically dominant cities (Rotterdam, Antwerp) as well as linguistically more peripheral cities (Breda, Dendermonde). More cities could have been included, but this would have complicated the design of the model by quite a margin, and would also have made the model more computationally intensive.

Every influence sphere includes an appropriate population of agents, which is based on historical demographic data by the Dutch Bureau of Statistics (Centraal Bureau voor de Statistiek 1930) and Hertogen (2013). The population censuses used are the ones from 1930[3] for both the Netherlands and Belgium, with a ratio of one agent for every 3000 people in the real historical populations. To keep the simulation as simple as possible, the population counts remain fixed — no agents are added or removed, and the distribution among cities remains the same for the remainder of the entire simulation. Should cities grow and shrink during the course of the simulation run, it would be very difficult to still assess the effect of the assumptions we are trying to test. For example, should one city grow much bigger during the simulation run, its influence will grow as well (cfr. infra), which could affect the results of a possible convergence or divergence. Such effects are currently avoided by ensuring that the cumulative population counts of the Netherlandic and Belgian spheres are comparable, which is also why Belgium is represented by one extra city in the simulation. As such, we know that any influence on the convergence-divergence situation stems from the current assumptions built into the model. Should we allow cities to grow independently of each other, this population equality between the two countries is no longer guaranteed, which will make the model results opaque and difficult to interpret.

2.2.2 Travel

Inhabitants of the aforementioned influence spheres are free to move around the model world. Movement in the simulation world is *always* restricted to one step into one of the four cardinal directions per unit of time[4] in the model. Agents thus have to move one step at a time. Diagonal movements are not allowed.

There are two 'movement modes' for agents in the simulation: wandering and travelling. When an agent has no specific travel destination, they simply wander. This means that they move one unit into one of the four allowed directions, or stay in place. At every step an agent is not travelling somewhere, the agent has a $\frac{1}{20}$ chance of moving back towards the centre of the home sphere. This means that they keep moving one step towards the home sphere centre per unit of time in the model until they are home. This process can be compared to a chess pawn which has to move in small

2. The model developed in this study is available online at github.com/AntheSevenants/BorderModel. The research data, including R scripts, is available online at github.com/AntheSevenants/BorderModelData.
3. Remember that this is the time period in which the Dutch pronunciation in the Netherlands started to shift.
4. Time in the computer model is not expressed in terms of minutes or hours, but rather in terms of abstract units. At every unit or step of time, all agents move and speak (in that order) according to predefined, simple logic. In theory, a simulation could run endlessly for an infinite number of these abstract units. In reality, every simulation will run for a set number of steps (see below).

increments across multiple turns in order to reach a destination. The return mechanism is put in place to make sure that agents do not wander off too far so that the notion of distinctive cities is lost, and is comparable to the homing mechanism found in Stanford and Kenny (2013), which serves the same purpose there.

In addition to simply wandering, agents can also purposely travel to another sphere. It is clear that agents should not visit each sphere equally as often, since in real life not every destination is equally popular either. At every step of time in the model, every agent has the possibility to enter a 'ready to travel' state.[5] The agent can enter this ready state either for domestic travel or for abroad travel, but not for both at the same time. The probabilities for domestic and abroad travel are controlled separately. Domestic travel is prioritised in the sense that the probability for domestic travel is tested first. If this probability is set to, for example, 0.005, an agent has a $\frac{1}{200}$ chance of entering the ready to travel state for domestic travel. If this probability is not met (in our example, this is a $\frac{199}{200}$ chance), the abroad travel probability is tested. If this probability is set to 0.001, an agent has a $\frac{1}{1000}$ chance of entering the abroad ready to travel state. If no travel is initiated (because both probabilities were not met), the agent simply wanders. The domestic and abroad travel probabilities (the $\frac{1}{200}$ and $\frac{1}{1000}$ chances in our example) can be controlled separately for each country. In order to make everything clearer, an overview of all four probabilities and their uses is given in Table 1.

Travel probability	**Use**
domestic travel probability (NL)	the probability governing how likely an agent from the Netherlands is to travel to another sphere in the Netherlands
abroad travel probability (NL)	the probability governing how likely an agent from the Netherlands is to travel to a sphere in Belgium
domestic travel probability (BE)	the probability governing how likely an agent from Belgium is to travel to another sphere in Belgium
abroad travel probability (BE)	the probability governing how likely an agent from Belgium is to travel to a sphere in the Netherlands

Table 1: An overview of the four general travel probabilities. There are actually only two general travel parameters, but they can be controlled separately for each country in the simulation. This makes for four parameters in total.

Once an agent reaches the ready to travel state, a random influence sphere is selected. If the agent is in a domestic ready to travel state, a location from the agent's own country is selected. If the agent is in an abroad ready to travel state, a location from the other country is selected. Whether an agent then actually travels to this destination is dependent on the probabilities derived from the gravity model (Trudgill 1974). Trudgill's gravity model attempts to model the linguistic influence of one city on another, taking into account both the distance between the cities and their respective population counts. One implementation of the model, which calculates the influence of city i on city j ($\mathrm{I_{ij}}$), is given in Trudgill (1974):

$$I_{ij} = s \cdot \frac{P_i P_j}{(d_{ij})^2} \cdot \frac{P_i}{P_i + P_j} \tag{1}$$

The first part of the equation, s, stands for an index of linguistic similarity. This is an index between 0 and 4 which one can assign to express the similarity between the language varieties of the influencing city and the influenced city. The higher the index, the more similar both varieties are. Trudgill's proposal for these indices is, however, rather unsystematic, and it is unclear what

5. When agents are already travelling or going back to their home sphere (see above), the travel possibility is temporarily disabled until they are no longer travelling or have returned home.

formal criteria should be used to decide which indices should be assigned to which cities. This unsystematicity in addition to the idea that the linguistic distance between standard languages should always be minimal led to the decision not to use linguistic similarity in the simulation. The next part of the equation, $\frac{P_i P_j}{(d_{ij})^2}$, describes the interaction between the two cities. The numerator of the fraction simply multiplies the populations of the influencing city (P_i) and the influenced city (P_j). The denominator is the squared distance between both cities $(d_{ij})^2$, which implies that the interaction between both cities weakens quadratically with distance. Finally, this *interaction* index is turned into an *influence* index ($\frac{P_i}{P_i+P_j}$). The interaction index is multiplied by the population of the influencing city (P_i) divided by the sum of the influencing city and the influenced city ($P_i + P_j$). This has the effect that larger cities will have a much larger impact on smaller cities than vice versa, even when their interaction indices are the same.

The output of the gravity model cannot be used as-is in the simulation, however, as it computes a metric of influence, not a probability of travel. To solve this problem, the different influences on a specific city are compared and then converted to a probability according to the Equation 2. This happens separately for every country:

$$p_{ji} = \frac{I_{ij}}{I_j} \tag{2}$$

p_{ji} is the probability of an agent from sphere j visiting sphere i, I_{ij} is the influence of i on j (as computed by the gravity model). I_j sums up the total influence on j (from one country). As one can infer from the equation, the travel probability of an agent of influence sphere j visiting influence sphere i is dependent on the influence of i on j, not j on i. This reversal is self-evident, as we can imagine that the influence of Antwerp speakers on Turnhout speakers stems from Turnhout's inhabitants visiting Antwerp, and not vice versa. Without this inverse relationship, the model produces nonsensical results. An example of the probabilities generated by the gravity model for the city Leuven is given in Figure 1.

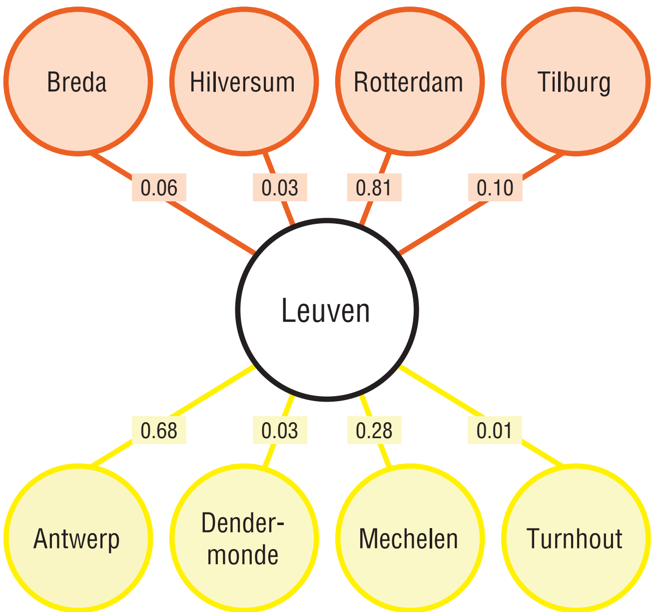

Figure 1: The exact travel probabilities for agents of the city Leuven as generated by the gravity model. Abroad travel destination probabilities are given above. Domestic travel destination probabilities are given below.

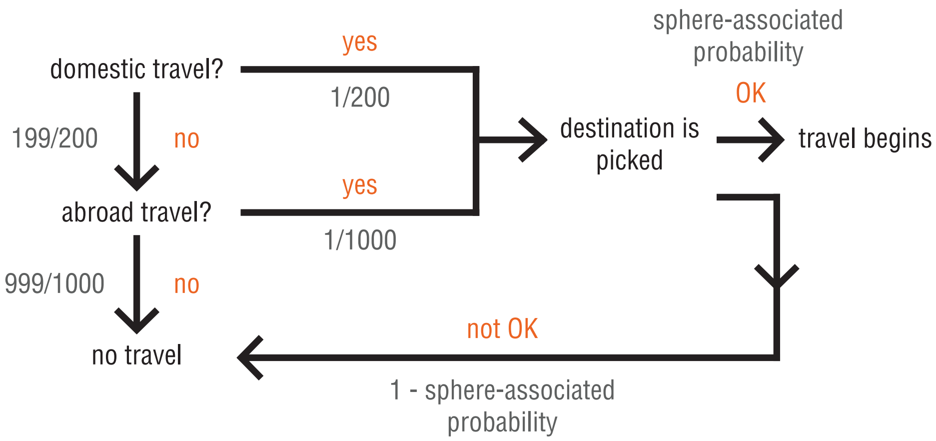

Figure 2: The travel logic represented as a flow chart. Example probabilities are given. The start of the chart is in the upper-left corner. The logic in this chart decides whether an agent can move from the wandering state to the travelling state in Figure 3 (both indicated in orange). If the logic ends in 'no travel', the state remains the same. This logic is evaluated at every step of the model, barring the conditions discussed in Footnote 5.

We saw that once an agent reaches the ready to travel state, a random influence sphere is selected. It is at this point that the probabilities generated by the gravity model come into play. An agent can only go to the selected travel destination depending on the probability associated with that sphere. Let us illustrate this for an agent from Leuven who has entered the abroad ready to travel state. The model has selected Tilburg as the agent's possible travel destination. We see in Figure 1 that for an agent from Leuven, Tilburg has a travel probability of 0.10 or $\frac{1}{10}$. The agent thus has a $\frac{1}{10}$ chance of actually leaving for Tilburg. If this $\frac{1}{10}$ probability is met, the trip can commence. If it is not met, the agent simply returns to the wandering state. This second layer of the travel destination selection, based on the gravity model, makes sure that the frequency at which agents visit other spheres is realistic. A flow chart of the complete travel logic is shown in Figure 2.

After an agent has been presented with the opportunity to travel and the trip is 'greenlit' by the gravity model mechanism, the agent moves in a straight line towards the centre of the destination's influence sphere. The movement towards this sphere is constrained by the same general restrictions as wandering (one step per unit of time in the model, only cardinal directions, no diagonals). An agent 'arrives' once they reach within half the radius of the destination sphere from its centre. They thus do not need to reach the absolute centre of the sphere. Once arrived, the agent starts wandering again, but at every step, they have a $\frac{1}{20}$ chance of returning home. The trip home also features the same travel and movement restrictions, including the half-radius rule. An obvious difference is of course that once an agent is back home, they enter the default 'wandering and homing' state again, until they go on another trip. An overview of all possible movement states is shown as a finite-state automaton in Figure 3.

2.2.3 Speaking

Of course, agents should also be able to talk to each other. After an agent has had the opportunity to move — be it as part of a path towards a travel destination or simply as the result of wandering — they are allowed to 'speak' to one 'hearing' agent in their neighbourhood, which consists of all

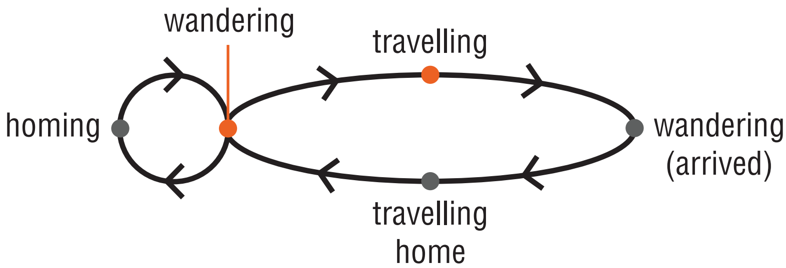

Figure 3: Agent movement logic in the model represented as a finite-state automaton. Every dot represents a movement state. Every arrow represents a state change. Agents can remain in a specific state for multiple time units. The wandering and travelling states are indicated in orange and correspond to the possible outcomes of the travel logic described in Figure 2.

neighbouring agents in the vicinity of eight cells of the speaker (so including diagonal cells[6]), and also agents occupying the same cell. From this neighbourhood, one hearer is randomly selected. Because of this random selection, conversations need not necessarily be reciprocal if multiple agents are close to each other. When an agent has no neighbours, they do not speak.

During a conversation, the speaker recovers a random sound from their sound repository, a collection of sounds which the speaker has heard before, and 'speaks' this sound to the hearer. The hearer then stores this sound in their own sound repository. Note that a sound repository can hold duplicate sound values; it is a collection of all recently heard sounds, not a collection of *kinds* of recently heard sounds. By speaking, a speaker does not influence their own sound repository; they for example do not duplicate the sound they spoke in their own repository. Changes in an agent's sound repository thus always come from another agent.

To simulate the effect of 'forgetting' older sounds, a sound repository holds a maximum of 140 forms, as in Stanford and Kenny (2013). When the sound repository is full, the oldest sound is removed after a new sound is added. Sounds are represented as real numbers between 0 and 1 as an abstraction of every possible Netherlandic innovation. Real numbers were chosen instead of binary values in order to be able to also generalise over vowel innovations, which traditionally do not follow binary oppositions. To guarantee that a shift towards the maximum value actually takes place, or, in other words, in order to ensure that the shifts described in Van de Velde (1996) happen, a speaker from the Netherlands will only add Netherlandic sounds to their sound repositories if they are higher than the current average value of their sound repository. The same speakers from the Netherlands can still be influenced by Belgian speakers and sounds. It might be argued that this behaviour makes the model 'circular' and thus designed in such a way that only the desired outcome is possible. One has to take into account, however, that the present simulation does not aim to uncover how the many shifts in the Netherlandic standard language described in Van de Velde (1996) spread. Instead, the shifts are implemented as a 'given' in the model, and all latent factors which caused this shift, whatever those may be and however many these may be, are assumed. As such, we can fully focus on the effect the shift has on Belgian speakers, which is the actual topic at hand.

Speakers from the Dutch Randstad area[7] are initialised with near-maximum values[8] in their sound repositories — the Netherlandic innovations are said to stem from here (van Hout et al. 1999),

6. The no diagonal movements restriction was put in place to make sure that each step in the model has an equal distance (diagonal movements would have a distance of $\sqrt{2}$ in the model). Since no movement is involved in this case, diagonals are allowed.
7. '[T]he area comprising the cities of Amsterdam, Rotterdam, the Hague and Utrecht' (Willemyns 2003). In the model, the cities of Rotterdam and Hilversum are located within this Randstad area.
8. The mean of the Randstad values is 0.89. Sounds are generated within an 0.1 interval around this mean, which means that the actual sound values range from 0.79 to 0.99.

while all other regions start with near-zero values.[9] Note that this does not mean that we define 140 particular linguistic contrasts in the Dutch language, nor does this mean that *all* sounds in the Dutch language differ. Rather, by filling the initial repositories of the Randstad agents with high sounds values, we initialise them with a general high degree of 'Netherlandic-ness'. If an agent's sound repository contains many values which are close to 1, then the idea is that this agent represents a speaker which generally opts for the innovative Netherlandic forms rather than the older conservative forms. If an agent's sound repository only contains low values, then the agent represents a speaker of the older language forms. Agents with both high and low values are possibly undergoing an evolution. The sound range is thus *only* applicable to all the possible conservative-innovative contrasts and makes an abstraction of those contrasts.

It would be theoretically possible to individually model the different linguistic contrasts over which the model currently abstracts. Moreover, it would be possible to also incorporate the sounds which do not have a contrastive difference. The problem with this approach is that it does not add much to the model design, since we are only interested in a general notion of 'degree of innovation'. Since the sounds which are not contrastive do not contribute to this degree of innovation, and since the model does not care about exact innovations, it does not make sense to build these aspects into the model.

Strictly speaking, the strong establishment of innovative forms in the Randstad area is ahistorical and goes against the course of the shifts described in Van de Velde (1996). In that study, it is clear that around the 1930s — the starting point for these simulations — most Netherlandic innovations were still in their infancy. In our model, the innovations are much better established. Unfortunately, this is a problem which cannot be solved within the current implementation of speaking in the model. Dutch agents cannot 'produce' new sounds themselves[10], so the end value of the evolution should already be present somewhere in an agent's sound repository. In order to guarantee that the sound evolution in the Netherlands actually carries out completely, the end value should be somewhat well-represented.

Every agent starts with a full, 140 sound inventory. Given the gentle slope with which typical diffusion curves, curves which describe the spread of a new linguistic innovation, start (Bailey 1973), agents would be easily and immediately influenced by new sounds they hear if they started with fewer sounds in their repositories, which would be unrealistic. A schematic representation of the sound exchange is given in Figure 4.

2.2.4 Testing the theories

Up until now, only the basic behaviour of the model has been explained. To test the research questions described above, we need to systematically change the parameters of the model and subsequently observe the model output. We test the first theory, which relates the divergence between the two standard language pronunciations to a lack of contact between Belgian and Dutch speakers, by varying the probability of agents travelling abroad. The tested probabilities range from absurdly low (agents not travelling abroad at all) to absurdly high (agents visiting cities in the other country more often than local cities), in order to cover the entire spectrum of possible scenarios. The domestic travel probabilities for both countries remain fixed at $5 \cdot 10^{-3}$.

The second theory, which predicts that the pace of language innovations in the Netherlands was too fast for Belgians to be able to follow them, is tested by varying the number of times a spoken sound is saved to a Dutch hearer's sound repository. Normally, every agent only saves a newly encountered sound in their sound repository once. When a new sound exemplar is added multiple times to an agent's sound repository, however, this increases the representation of that sound in that agent's sound repository, which also increases the pace at which new sounds are adopted in the

9. Precisely 0.00001.
10. This limitation is the result of a design choice in this specific model, and is not a deficiency of simulation research as a whole. It would be possible to replace the current copying implementation for speaking with another one, which could then for example hinge on analogy or random variation.

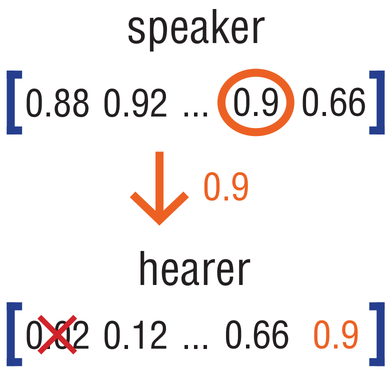

Figure 4: A schematic overview of a sound exchange between speaker and hearer. The blue brackets represent the sound repositories of the agents. The numbers in these brackets are the sound values stored in the repositories. In this example, the circled sound is randomly selected from the speaker's sound repository and is subsequently saved in the hearer's sound repository. Because an agent's sound repository is initialised with 140 sounds, every addition to the sound repository guarantees a deletion. Therefore, the leftmost value in the repository (assumed to be the oldest value) is deleted.

Netherlands. The number of exemplars saved is varied from once to up to 20 times. The abroad travel probabilities for both countries remain fixed at 10^{-3}, the domestic travel probabilities at $5 \cdot 10^{-3}$.

The third theory, which claims the divergence between the standard language pronunciations is the result of a Belgian aversion of Netherlandic Dutch, is a little more complex in the simulation. In the model, agents receive a value for 'ethnocentrism'. The higher this value is, the less likely an agent is to store a sound from another country in their sound repository. For Dutch speakers, this property is always set to 0.85 in order to reflect the asymmetric language hierarchy between the Netherlands and Belgium. Belgian speakers have varying ethnocentrism values across simulation runs. The simulation runs start at no ethnocentrism for all Belgian speakers, which means all foreign sounds will be adopted. Then, every next simulation run, the ethnocentrism value will be collectively increased by 0.01 until maximum ethnocentrism is reached in Belgian speakers, which means no Belgian agent will adopt any more foreign sounds. The collective increase means that every Belgian agent indeed shares the same ethnocentrism value within one simulation run. We test the entire range of values in order to get a complete picture of the influence of ethnocentrism on the adoption of foreign innovations.

In addition, a geographical model of ethnocentrism is also tested. This geographical model relates the values for ethnocentrism for agents in both countries to how close a specific agent lives to the border. Based on the idea that speakers who live on a border with another country are more interested in what happens in that country (NOS Afdeling kijk- en luisteronderzoek and BRT Studiedienst 1983), the closer an agent lives to the border, the lower their ethnocentrism value will be; agents who live on the border receive a zero value, agents who live a maximal distance away from the border receive the maximum value, one. In this geographical model, speakers from the Netherlands are also assigned variable values, unlike in the fixed model. The geographical ethnocentrism model is used as an extension of the model with fixed ethnocentrism values for the entire population. The abroad travel probabilities for both countries remain fixed at 10^{-3}, the domestic travel probabilities at $5 \cdot 10^{-3}$.

The final theory, which relates the linguistic evolutions in both countries to the media landscape in those countries, is tested by simulating media influence. When media influence is activated, Dutch agents 'receive' sounds from the Randstad area at every step in addition to their regular conversations. This is the linguistically dominant region in the Dutch media. Belgian agents receive sounds from the Brabant area, which is the linguistically dominant area in the media in Flanders. In addition, Belgian agents can also receive the same Dutch Randstad sounds, since Dutch television was very popular in Flanders in the past (Geeraerts 2017). The ratio is $\frac{1}{4}$ for Dutch sounds, and $\frac{3}{4}$ for Belgian sounds. These ratios are based on a report by the Flemish National Television (Instituut der Nederlandse Uitzendingen 1982). Sounds from the central areas are sampled by selecting a random agent from the specific central region (Randstad or Brabant), and then selecting a random sound from that agent's sound repository. This happens independently for every agent, so each agent receives different media sounds.

Media receptiveness is defined individually for each agent separately, since receptiveness to media innovations seems to be person-bound (Stuart-Smith and Timmins 2009). Whether a media sound is actually saved to an agent's sound repository depends on the media receptiveness of that agent. The higher the receptiveness, the higher the adoption probability will be. Even though media receptiveness is defined individually, it is varied across the entire population systematically — the entire population receives the same value in one simulation run. This is done in order to keep the simulation as straightforward as possible. The abroad travel probabilities for both countries remain fixed at 10^{-3}, the domestic travel probabilities at $5 \cdot 10^{-3}$.

2.3 Results and discussion

Every simulation theory was tested by varying the different parameters associated with the respective theories according to the principles outlined above. Every simulation with a specific set of parameters was run for 25,000 steps. We assume this range of 25,000 steps to correspond to the period spanning the 1930s until now. In principle, we could have chosen any other 'sensible' number (with 'sensible' meaning that we do not run the simulation for just ten steps, for example), but this relatively high number was chosen specifically to give parameter sets conducive to very slow diffusion a chance to complete this diffusion. All simulations were run on a powerful virtual machine generously provided by Mads Janszen. We are very grateful for his support.

The simulation output, which includes — among other data — the degree of sound innovation for every region and country, was analysed by interpreting the sound evolutions directly.

2.3.1 Contact theory

The contact theory posits that reduced contact between speakers from the Netherlands and Belgium will result in a divergence between the standard pronunciations in both countries. We see in Figure 5 that reduced contact can indeed lead to divergence[11] in the model, but only starting from an abroad travel probability of 10^{-6}, a travel rate 5,000 times smaller than the domestic travel rate, which is fixed at 0.005. Travel rates higher than 10^{-6} are always conducive to some degree of convergence.[12]

The interpretation of the ratio between the default domestic travel parameter (0.005) and the abroad travel probability found is complex. While the domestic travel probability is many magnitudes bigger than the abroad travel probability, the imbalance is not completely unrealistic. For many inhabitants of both countries, even those who live close to the border and are willing to travel to the other country, life primarily takes place within the borders (as evidenced by the attitudinal questions in NOS Afdeling kijk- en luisteronderzoek and BRT Studiedienst (1983)). Even if people occasionally go on trips to the other country, these trips will be counterbalanced by many more in-

11. In this context, divergence has to be interpreted as a situation in which Belgian speakers clearly do not follow northern innovations. Every situation which shows even a slight tendency to follow the innovations will be regarded as convergence, since it follows the innovative trend in the Netherlands (which is not shown here).

12. If we had chosen to end the simulations at, for example, 2,500 steps, 10^{-5} and possibly even 10^{-4} would have become viable divergence parameters. The choice for 10^{-6} is thus entirely conditioned by our parameter settings, but the result, a large discrepancy between the abroad and domestic travel rates, remains the same. The exact order of magnitude does not play a role in the interpretation of this discrepancy either.

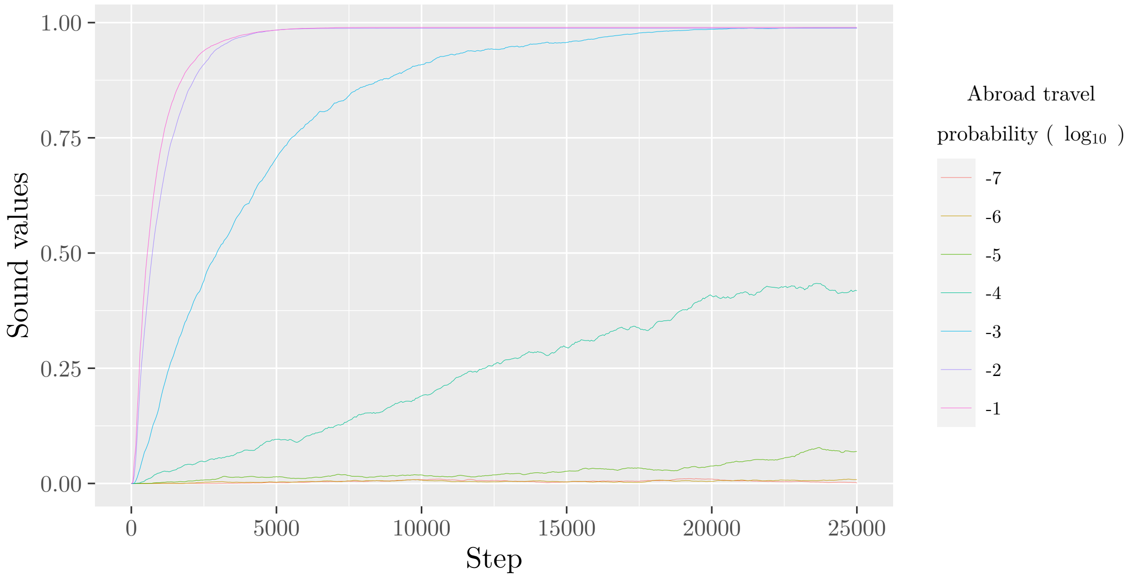

Figure 5: A visualisation of the sound evolutions in Belgium in the simulation model. Every colour represents a simulation run with a different value for the abroad travel probability parameter in Belgium and the Netherlands. For clarity's sake, the probabilities are shown as $\log_{10}$.

teractions in the home country. In this light, the ratio seems more reasonable. Of course, the model does not contain the more intricate details of travel, such as highways, natural borders or travel cost. Though it is possible to implement these properties, more complex travel implementations do not always yield better results (Nerbonne et al. 2005). It seems then that the notion of contact abstracts quite straightforwardly to a simulation context, which means that our contact model should be a satisfactory approximation of reality. While one must be cautious not to take the probabilities and ratios presented here at face value, our implementation of the contact theory in the model produces the results outlined by the theory. In our computer model, then, the proof of concept of this theory indeed works.

2.3.2 Moving target theory

The moving target theory predicted that an increased pace of language change in the Netherlands could cause a divergence in Belgium, as the result of language evolving 'too fast'. In Figure 6, we see that for most of the course of the graph, the difference between the various target acceleration values seems negligible. It is thus doubtful whether target acceleration has any impact on convergence or divergence at all.

This does not mean, however, that the theory as a whole should be written off. One reason is that we only checked in the computer model whether just the pace of language evolution is enough to cause divergence in Belgium. Another part of the moving target theory, which is much harder to capture in a simulation, is the presumed exasperation experienced by Belgian speakers who cannot keep up with Netherlandic innovations. Incorporating a form of 'cut-off' in the model, where Belgian speakers would refuse to adopt Netherlandic sounds if they had shifted too fast or too far, would make the model circular and diminish its value. Still, there is no reason to believe that such processes do not exist in the real world, but they simply make no sense as part of an explorative simulation model like this one. Nevertheless, the implementation of this theory without the presumed Belgian exasperation does not work as a proof of concept in our model. The model does not produce the divergence the theory prescribes. The results of the moving target theory do tell us that in our proof

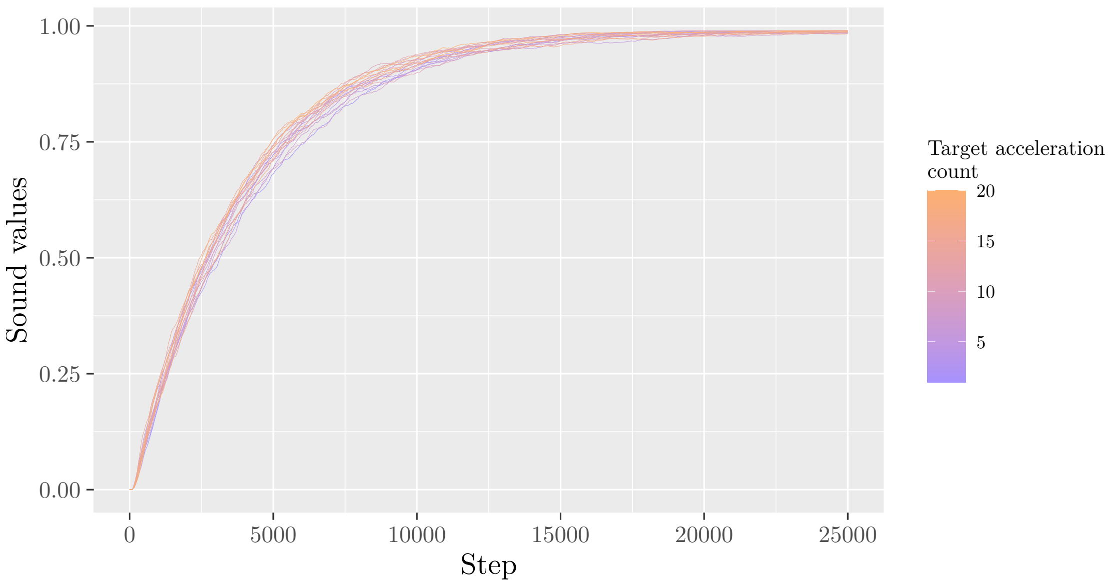

Figure 6: A visualisation of the sound evolutions in Belgium in the simulation model. The colour scale represents the different values for the target acceleration parameter in the Netherlands.

of concept, given a considerable amount of contact between two countries (abroad travel rate is $\frac{1}{5}$ of the domestic travel rate), and without exasperation, pace has virtually no influence. It would be interesting to examine this relationship in more detail in further research.

Another reason why the theory as a whole should not be written off stems from the contact theory. Indeed, one could argue that the moving target theory and the contact theory are actually two sides of the same coin. The idea behind this is that the effect of a lack of contact between two countries (contact theory) also depends on the pace of the language evolution in the country which one country is 'meant' to follow (moving target theory). It seems logical that a lack of contact is only a problem when this means that the dynamic in that country can no longer be sufficiently 'picked up' by people in the other country. Simply put, if the pace of language change in the 'leader' country speeds up, the contact rate of the 'follower' country should follow along to compensate for this. This is a dynamic which was not tested in the isolated implementations of both the contact and the moving target theories. The results of the moving target theory do tell us that, given a considerable amount of contact between two countries (abroad travel rate is $\frac{1}{5}$ of the domestic travel rate) and without exasperation, pace has virtually no influence.

2.3.3 Ethnocentrism theory

The ethnocentrism theory posits that more ethnocentrism in Belgian speakers leads to less adoption of Netherlandic sounds. Figure 7 shows that the higher the ethnocentrism values are in the simulation, the slower the pace of adoption is. In addition, we also see that ethnocentrism values at the higher end of the scale can lead to divergence in Belgium in the model. For this specific theory, however, the abroad travel rate had to be lowered from the default 10^{-3} to 10^{-4}. With the abroad travel rate fixed at 10^{-3}, the model was only conducive to convergence, as abroad travel seems to be a 'stronger' parameter in this model than ethnocentrism. Since we know from the contact theory results that abroad travel probabilities by themselves only cause a divergence starting from 10^{-6}, if we see a form of divergence at a travel rate of 10^{-4}, we know that it must stem from the ethnocentrism system.

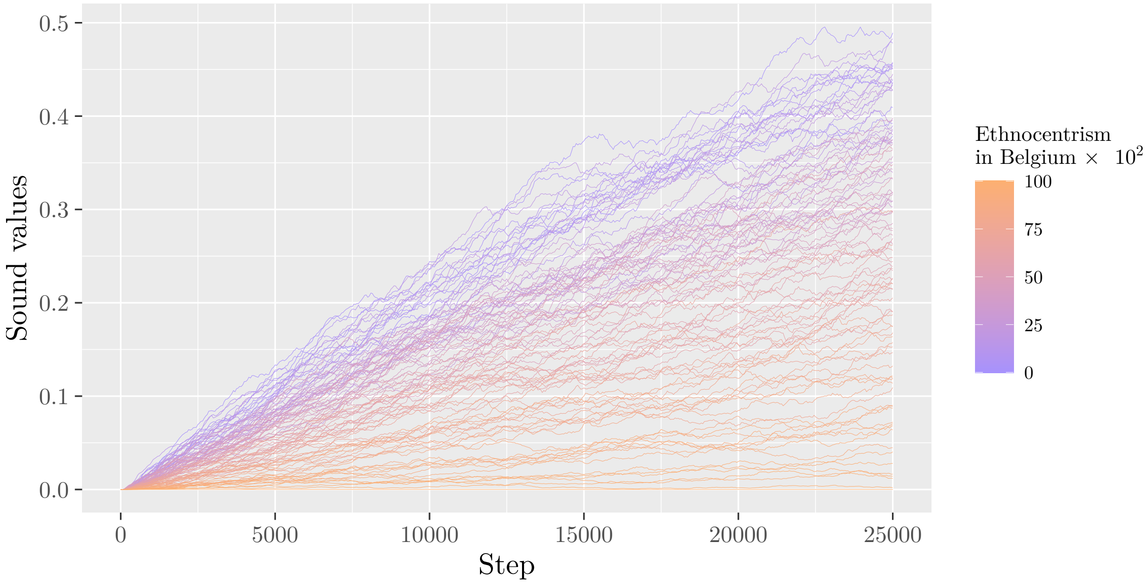

Figure 7: A visualisation of the sound evolutions in Belgium in the simulation model. The colour scale represents the different values for the ethnocentrism parameter in Belgian speakers. For clarity's sake, the ethnocentrism values are shown multiplied by 10^2. The abroad travel rate in this set of simulations is fixed at 10^{-4}.

The results of the scaled ethnocentrism theory are comparable to simulation runs with low to no ethnocentrism in Belgian agents. It thus seems that the way ethnocentrism values are distributed across the agent population has an impact on what effects it has on the simulation outcome. A possible explanation for this behaviour portrays border agents as gatekeepers. Since border agents have much lower ethnocentrism values, they are more inclined to adopt foreign innovations, making them their own. As such, these foreign innovations now belong to indigenous sound repositories, and, as a consequence, can now be adopted by high ethnocentrism agents who no longer deem the sounds foreign.

2.3.4 Media theory

The final theory, the media theory, dictates that increased media influence will amplify existing tendencies for language change, be it acceleration or inhibition. We see in Figure 8 that media receptiveness can accelerate convergence to an impressive extent. The bottommost line shows the baseline situation without any media influence applied. We see that, as the media receptiveness values increase, so does the pace at which convergence happens. This means that, if the linguistic climate in Belgium is conducive to convergence, media influence will accelerate this convergence, at least in our model.

At the same time, we see here that we cannot use these simulations for checking how the media theory interacts with a divergence base situation — the situation without any media receptiveness is already one of convergence. Therefore, another batch of simulations had to be run. Luckily, we know from the results of the contact theory that no convergence occurs when the abroad travel parameters are set to 10^{-6}. As such, if we start from this divergence situation with an abroad travel probability of 10^{-6} and add media influence to it, we know that any changes in the behaviour of the sound curvature are due to this added media influence. This alternation is only meant to investigate the possibilities of the theories within the model space, not to prove any relationship between them. The results of the second batch of simulations are shown in Figure 9. Surprisingly, we see that

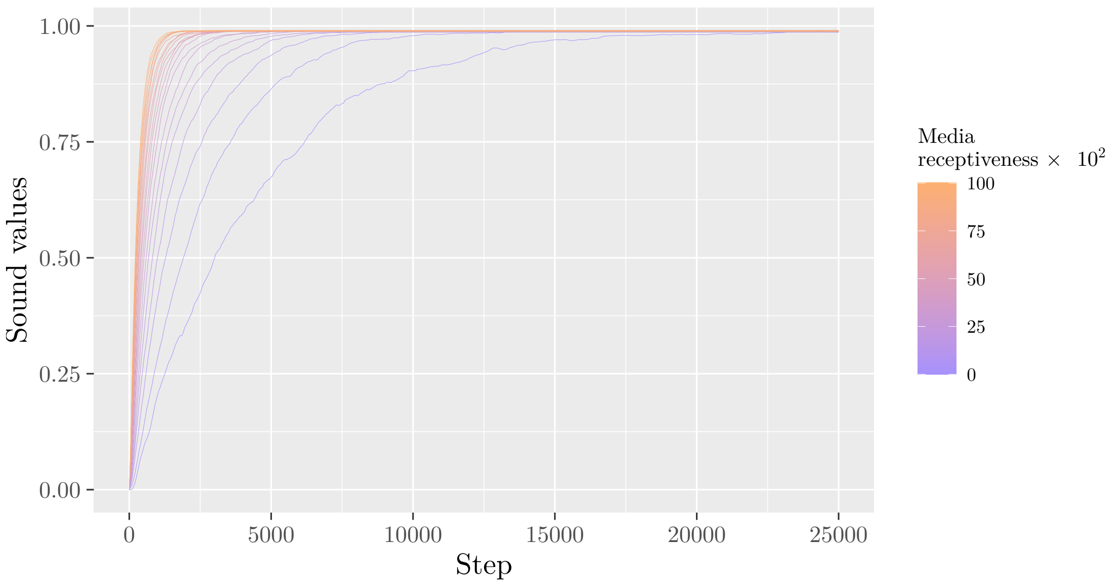

Figure 8: A visualisation of the sound evolutions in Belgium in the simulation model. The colour scale represents the different values for the media receptiveness parameter in Belgian and Dutch speakers. For clarity's sake, the media receptiveness values are shown multiplied by 10^2.

the guaranteed divergence situation, represented by the zero media influence line at the bottom of the plot, immediately turns into a convergence situation when media influence is added. Even the smallest amount of media influence on the agents (0.01) causes convergence in the model.

It seems then that the expected inhibitory effect of the Belgian media influence does not manifest itself in the model. It is likely that the share of Dutch media influence ($\frac{1}{4}$) is high enough to kickstart convergence at even the lowest parameter settings, despite the higher share of Belgian media influence ($\frac{3}{4}$). The fact that media influence was implemented as a powerful force in the model (media influence has direct access to the sound inventory of all agents of an entire country – even two countries for Dutch media) could mean that 0.01 is still very influential. This does not mean that the media were equally powerful in real life – we know very little about the exact effects of media influence. Therefore, we have to interpret the model results under the assumption that media influence can have a large effect on language. If and only if we assume that media influence is a very powerful force, the media theory works in convergence situations in our model (as shown in Figure 8), but for divergence situations, the proof of concept deviates considerably from what the theory prescribes.

2.4 Shortcomings of the current model

In this section, we will evaluate the different aspects of the model and consider how they could be improved. The current implementation of space, using a grid and influence spheres, works well for simulation purposes. The grid system simplifies movement in the model and its cells make notions such as neighbourhood or adjacency immediately obvious. Furthermore, the influence spheres from Stanford and Kenny (2013) are also a good addition to the spatial dimension of the model. Their circular nature circumvents the problems that would result from attempting to capture the real-life shapes of cities and towns and their clear centre point makes mechanics such as homing or travelling easier to implement. The historical populations which inhabit these influence spheres and the distances between the spheres seem to be useful additions which make the model more realistic, but their potential may have been held back. The reason for this stems from the way sounds are initialised in the Netherlands at the start of the simulation. In this model, all Randstad influence

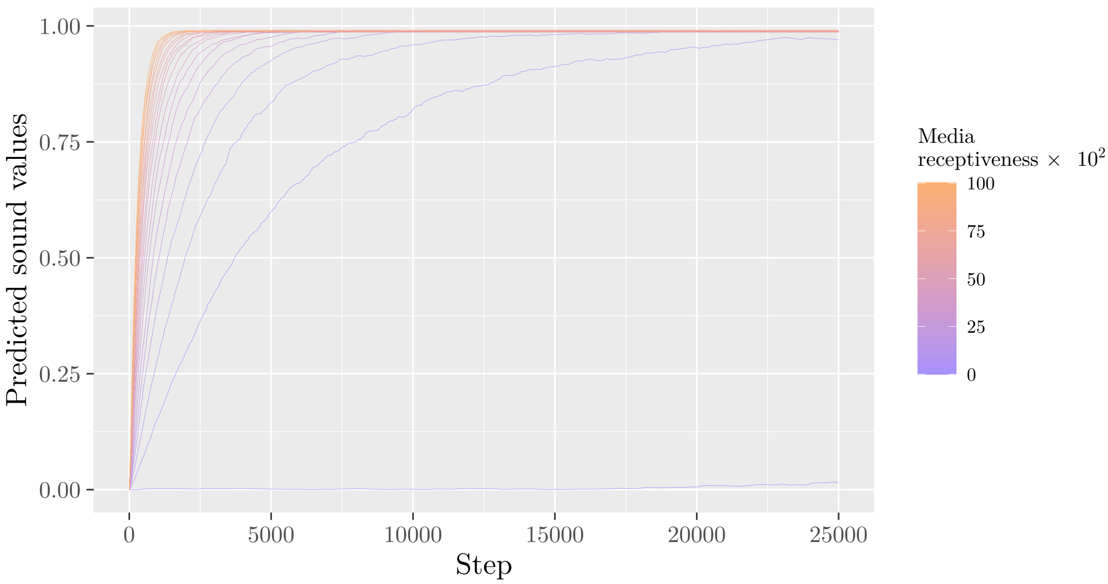

Figure 9: A visualisation of the sound evolutions in Belgium in the simulation model, second batch. The colour scale represents the different values for the media receptiveness parameter in Belgium and the Netherlands. For clarity's sake, the media receptiveness values are shown multiplied by 10^2. The abroad travel rate in this set of simulations is fixed at 10^{-6}.

spheres had to be initialised with high sound values, since the sounds associated with the end goal of the Netherlandic evolutions already had to be present from the beginning. This is the result of the fact that agents cannot create new sounds on their own. The problem with this sacrifice is that it raises the average sound values in the Netherlands to an unrealistic point (cfr. supra) and makes it so that the actual divergence is limited to only two spheres. Therefore, it is questionable whether in the current implementation the geographical component of the model had any meaningful contribution at all. At least, geography still had an influence on the pace of convergence, as is evident from Figure 10. In the graph, we see that spheres closer to the border (Antwerp, Breda) converged faster than spheres further from the border (Dendermonde, Leuven). Of course, geography was also important in the scaled ethnocentrism simulations, where the location of the different spheres played an important role (cfr. supra).

A solution to this 'high initialisation' problem would be to use hardcoded sound evolution trajectories for the Netherlands. Instead of recovering sounds from individual sound repositories, agents could also sample spoken sounds from a variable diffusion curve, depending on what step of the simulation we are currently in. Figure 11 shows how this could work, exemplifying the mechanism at step 300 of a fictional simulation. On the left side, we see a predefined sound evolution for the Netherlands. On the right side, we see a normal distribution centred around the value the sound evolution dictates for step 300. Every Dutch agent could sample from this distribution at step 300, from another distribution centred around the value for step 301 at step 301, and so on. For Belgian agents, the system with degrading sound repositories which is currently used should remain in place since it works well and is not plagued by the same issues the Dutch agents had. Of course, the curvature of the predefined sound evolution (on the left) could be varied to speed up or slow down the sound evolution in the Netherlands, which is especially interesting for the contact and moving target theories. The main advantage of this approach is that we can skip the Complex-Adaptive System entirely and instead focus on manipulating the 'result' of the system, without having to deal with all the problems it presents. It would no longer be necessary to have Dutch agents only adopt sounds

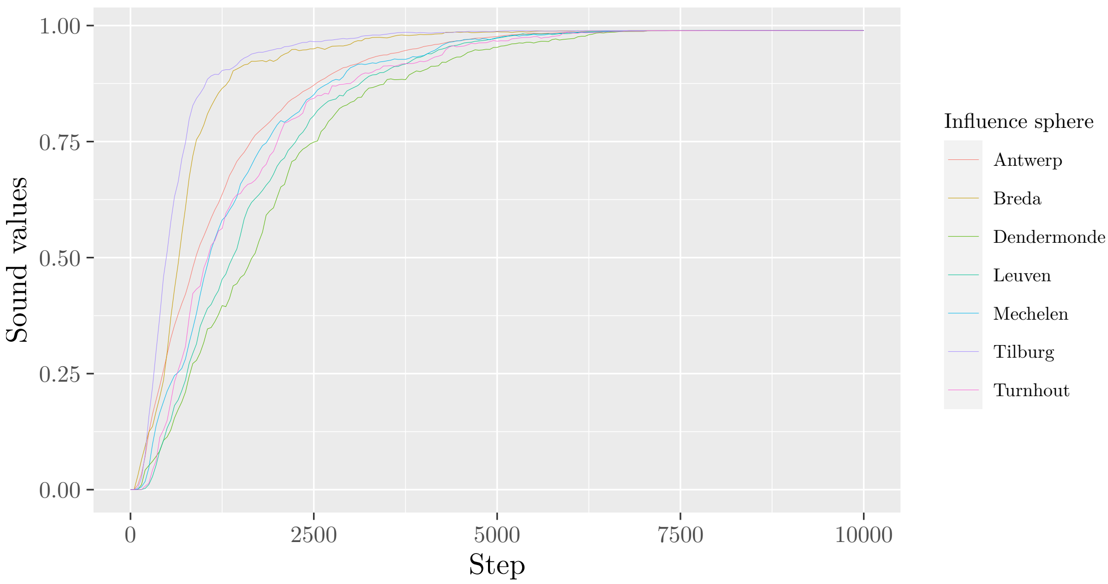

Figure 10: A visualisation of the actual sound evolutions in the contact simulation model. Abroad travel probability is set to 0.005 in this specific run. The colours represent the different cities in the simulation run. Rotterdam and Hilversum are excluded because of their high starting values.

which are higher than the current average sound of their sound repositories, for example, since they simply no longer have sound repositories. At the same time, skipping the complex adaptive system is the biggest disadvantage as well. We saw in the section about the modelling of speaking that the innovative Dutch sounds are said to originate from the Randstad area. If we use a globally controlled distribution for the entirety of the Netherlands, we lose the notion of the Randstad as a central area of innovation. Of course, we could make non-Randstad areas use a different curve which lags behind the Randstad area curve a little to simulate the idea of the Randstad areas taking the lead in the sound evolutions. Using hardcoded sound evolutions would make the divergence situation much more realistic, and would heighten the importance of the geographical location of the different influence spheres.

An element of the simulation which is closely related to the spatial component is travel. The wandering and homing systems which are currently used work well and achieve their goal of providing a sense of community for a specific sphere while still allowing the agents of that sphere to roam around more or less freely. Having the agents come back at random intervals kept the model world from falling into chaos, which would have defeated the purpose of having spheres represent different cities. The mechanism which regulates purposeful domestic and abroad travel, however, has room for improvement. As it stands, the travel probabilities for domestic and abroad travel are separate. It would be more straightforward if travel were only regulated by just one probability and the split between domestic and abroad travel governed by a simple proportion (e.g. 99% of travel should be domestic, 1% should be abroad). Moreover, the procedure which selects the travel influence sphere for an agent is needlessly complicated and deserves to be streamlined. Currently, the domestic and abroad travel probabilities do not actually represent how often agents travel. Rather, they represent how likely an agent is to enter the 'ready to travel state', or how often they are given the *opportunity* to travel. This is the result of the fact that the gravity model does not come into play until after an influence sphere is chosen. It would be much more transparent if every successful travel attempt would actually lead to travel, which would be possible if the probabilities generated by the gravity model influenced how often a certain travel sphere would be selected, instead of

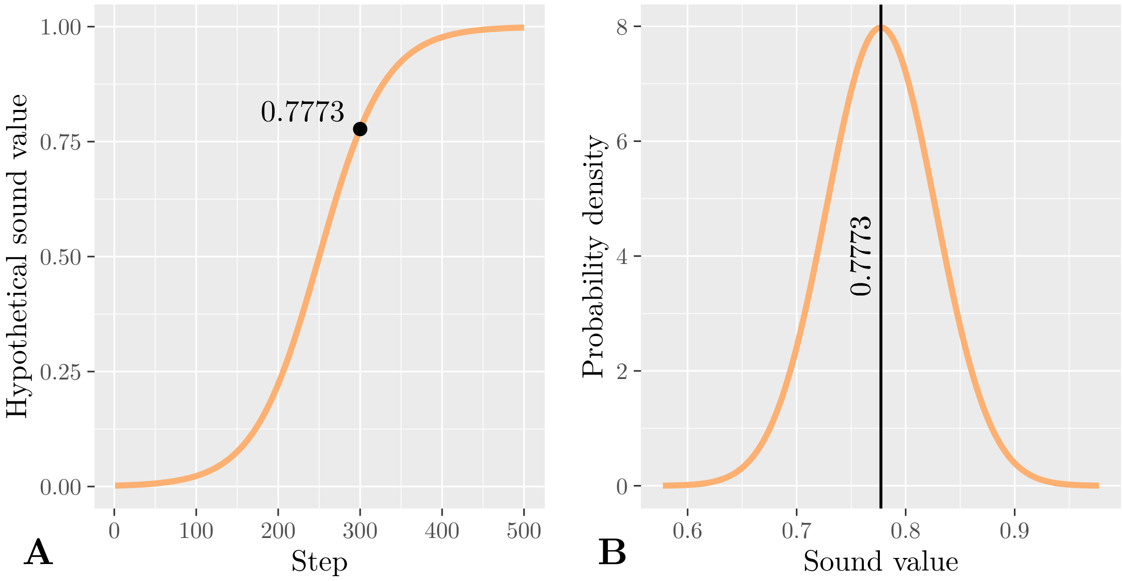

Figure 11: A fictional hardcoded sound evolution for the Netherlands. The hardcoded value for step 300 is highlighted. B: A density plot for a normal distribution centred around the value for step 300 from which sounds for Dutch agents could be sampled at that step.

the opposite implementation now. Of course, this means that to achieve the same travel rates of the current model, the travel probabilities would have to be adjusted downward, since in the new system, every opportunity to travel would actually lead to travel. Nonetheless, a streamlined travel selection procedure would make the interpretation of the probabilities much more straightforward, which would be useful during the interpretation of the model results.

3. Conclusion

In this article, we tested four theories presumed by the literature to be (at least partially) responsible for the divergence of the Dutch and Belgian standard languages by using agent-based computer simulations. The results of the simulations show that a lack of contact between both countries can lead to divergence in the model, but only if abroad travel is considerably less likely than domestic travel. The pace of language change in the Netherlands does not have a sizeable impact on convergence or divergence tendencies in Belgium in our proof of concept. High values for ethnocentrism in Belgian agents are able to lead to divergence in the model, as long as these high values are shared by the entire population. If ethnocentrism decreases along with how close agents live to the border, it has little effect. Media receptiveness in agents always kickstarts convergence in the model and it accelerates this convergence as well. Since media influence is implemented as a powerful force in the simulation, this result must be interpreted from the viewpoint of media having a sizeable impact on language change.

4. Acknowledgements

We would like to thank Freek Van de Velde and Dirk Pijpops for finding the time to discuss the architecture of the simulation model and giving their invaluable feedback.

References

Bailey, Charles-James N. (1973), *Variation and linguistic theory*, Center for applied linguistics, Arlington.

Baronchelli, Andrea (2014), Modeling is a tool, and data are crucial: A comment on "Modelling language evolution: Examples and predictions" by Tao Gong et al., *Physics of Life Reviews* **11** (2), pp. 317–318. http://www.sciencedirect.com/science/article/pii/S1571064514000153.

Beuls, Katrien and Luc Steels (2013), Agent-Based Models of Strategies for the Emergence and Evolution of Grammatical Agreement, *PloS one* **8** (3), pp. e58960–e58960. Place: United States Publisher: Public Library of Science PLoS, Public Library of Science.

Beuls, Katrien and Remi van Trijp (2016), Computational construction grammar and constructional change, *Belgian Journal of Linguistics* **30** (1), pp. 1–13. Publisher: John Benjamins. https://www.jbe-platform.com/content/journals/10.1075/bjl.30.01beu.

Centraal Bureau voor de Statistiek (1930), Bevolking en bevolkingsdichtheid der gemeenten van Nederland op 1 januari 1930. Publisher: DBNL. https://www.dbnl.org/tekst/_bev001bevo07_01/.

Deprez, Kas (1985), De aard van het Nederlands in Vlaanderen, *Heibel: onafhankelijk kritisch-satirisch-polemisch tijdschrift* **19** (4), pp. 101–127.

Geeraerts, Dirk (2017), Het kegelspel der taal. De naoorlogse evolutie van de Standaardnederlandsen, *De vele gezichten van het Nederlands in Vlaanderen: een inleiding tot de variatietaalkunde*, 1 ed., Acco, Leuven.

Gong, Tao and Lan Shuai (2013), Computer simulation as a scientific approach in evolutionary linguistics, *Language Sciences* **40**, pp. 12–23. https://linkinghub.elsevier.com/retrieve/pii/S0388000113000600.

Gong, Tao, Lan Shuai, and Menghan Zhang (2014), Modelling language evolution: Examples and predictions, *Physics of Life Reviews* **11** (2), pp. 280–302. http://www.sciencedirect.com/science/article/pii/S1571064513001863.

Hertogen, Jan (2013), Bevolkingsevolutie 1831-2012 in gemeenten, arr., prov. en gewesten. Database. http://www.npdata.be/BuG/190-Bevolking-gemeente-1831-2012/Gemeente-1831-2012.htm.

Instituut der Nederlandse Uitzendingen (1982), BRT Jaaroverzicht 1982.

Kazil, Jackie, David Masad, and Nathan Vērzemnieks (2019), Mesa. https://github.com/projectmesa/mesa.

Kirby, Simon (1999), *Function, selection, and innateness: the emergence of language universals*, Oxford linguistics, university press, Oxford.

Kretzschmar, William A., Ilkka Juuso, and C. Thomas Bailey (2014), Computer Simulation of Dialect Feature Diffusion, *Journal of Linguistic Geography* **2** (1), pp. 41–57. Publisher: Cambridge University Press. http://www.cambridge.org/core/journals/journal-of-linguistic-geography/article/computer-simulation-of-dialect-feature-diffusion/5FC8D58C781597E1FAB72F2413EE1AAE.

Nerbonne, John, Ilse van Gemert, and Wilbert Heeringa (2005), A Dialectometric View of Linguistic "Gravity".

NOS Afdeling kijk- en luisteronderzoek and BRT Studiedienst (1983), *Media zonder grenzen?: het mediagebruik in de grensstreek van Belgisch en Nederlands Limburg*, sn, S.l.

Pierrehumbert, Janet B., Forrest Stonedahl, and Robert Daland (2014), A model of grassroots changes in linguistic systems, *arXiv:1408.1985 [nlin, physics:physics]*. arXiv: 1408.1985. http://arxiv.org/abs/1408.1985.

Pijpops, Dirk and Katrien Beuls (2015), Agent-gebaseerde modellering in de historische taalkunde. Een model van regularisatiedruk op de Nederlandse werkwoorden, *Handelingen der Koninklijke Zuid-Nederlandse Maatschappij voor Taal- en Letterkunde en Geschiedenis* **69**, pp. 5–23. Publisher: Koninklijke Zuid-Nederlandse Maatschappij voor Taal-, Letterkunde en Geschiedenis. https://lirias.kuleuven.be/retrieve/349212.

Pijpops, Dirk, Katrien Beuls, and Freek Van de Velde (2015), The rise of the verbal weak inflection in Germanic. An agent-based model, *Computational Linguistics in the Netherlands Journal* **5**, pp. 81–102. Publisher: Computational Linguistics in the Netherlands. https://lirias.kuleuven.be/retrieve/344980.

Smith, Andrew D. M. (2014), Models of language evolution and change, *WIREs Cognitive Science* **5** (3), pp. 281–293. https://onlinelibrary.wiley.com/doi/abs/10.1002/wcs.1285.

Stanford, James N. and Laurence A. Kenny (2013), Revisiting transmission and diffusion: An agent-based model of vowel chain shifts across large communities, *Language Variation and Change* **25** (2), pp. 119–153. http://www.cambridge.org/core/journals/language-variation-and-change/article/revisiting-transmission-and-diffusion-an-agentbased-model-of-vowel-chain-shifts-across-large-communities/4F504E6441AB67001CE0E1051517644D.

Stuart-Smith, Jane and Claire Timmins (2009), The Role of the Individual in Language Variation and Change, *in* Llamas, Carmen, editor, *Language and Identities*, Edinburgh University Press, pp. 39–54.

Trudgill, Peter (1974), Linguistic change and diffusion: description and explanation in sociolinguistic dialect geography, *Language in Society* **3** (2), pp. 215–246. https://www.cambridge.org/core/product/identifier/S0047404500004358/type/journal_article.

Van de Velde, Freek (2019), External history of Dutch. University lecture.

Van de Velde, Hans (1996), *Variatie en verandering in het gesproken Standaard-Nederlands (1935-1993)*. https://repository.ubn.ru.nl/handle/2066/146159.

Van de Velde, Hans, Mikhail Kissine, Evie Tops, Sander van der Harst, and Roeland van Hout (2010), Will Dutch become Flemish? Autonomous developments in Belgian Dutch, **29** (3-4), pp. 385–416. https://www.degruyter.com/document/doi/10.1515/mult.2010.019/html.

van den Toorn, Maarten Cornelis (1997), *Geschiedenis van de Nederlandse taal*, Amsterdam University Press, Amsterdam.

van der Sijs, Nicoline and Roland Willemyns (2009), *Het verhaal van het Nederlands: een geschiedenis van twaalf eeuwen*, Bakker, Amsterdam.

van Hout, Roeland, Georges De Schutter, Erika De Crom, Wendy J Huinck, Hanne Kloots, and Hans Van de Velde (1999), De uitspraak van het Standaard-Nederlands. Variatie en varianten in Vlaanderen en Nederland, *Artikelen Van de Derde Sociolinguïstische Conferentie*.

van Istendael, Geert (2005), *Het Belgisch labyrint: wakker worden in een ander land*, 16 ed., Arbeiderspers, Amsterdam.

van Oostendorp, Marc (2016), Engels praten met Vlamingen. https://www.youtube.com/watch?v=O8yukkUqf7Y.

Willemyns, Roland (2003), Dutch, *in* Deumert, Ana and Wim Vandenbussche, editors, *Germanic standardizations: past to present*, Impact: studies in language and society 18, Benjamins, Amsterdam, pp. 93–125.