Perception of linguistic rhythm by newborn infants

Franck Ramus
Laboratoire de Sciences Cognitives et Psycholinguistique (EHESS/CNRS), Paris, France

Previous studies have shown that newborn infants are able to discriminate between certain languages, and it has been suggested that they do so by categorizing varieties of speech rhythm. However, in order to confirm this hypothesis, it is necessary to show that language discrimination is still performed by newborns when all speech cues other than rhythm are removed. Here, we conducted a series of experiments assessing discrimination between Dutch and Japanese by newborn infants, using a speech resynthesis technique to progressively degrade non-rhythmical properties of the sentences. When the stimuli are resynthesized using identical phonemes and artificial intonation contours for the two languages, thereby preserving only their rhythmic structure, newborns are still able to discriminate the languages. We conclude that new-borns are able to classify languages according to their type of rhythm, and that this ability may help them bootstrap other phonological properties of their native language.
Key-words: newborn speech perception language discrimination rhythm prosody bootstrapping.

Human newborns, as young as a few days old, have intriguing speech perception capacities. For instance, they perceive phonetic contrasts categorically (Eimas, Siqueland, Jusczyk, & Vigorito, 1971), and they perceive well-formed syllables as units (Bertoncini & Mehler, 1981; Bijeljac-Babic, Bertoncini, & Mehler, 1993; Bertoncini, Floccia, Nazzi, & Mehler, 1995; van Ooyen, Bertoncini, Sansavini, & Mehler, 1997). In addition, it has been repeatedly suggested that they process linguistic rhythm, as revealed by their capacity to discriminate between different languages (Mehler et al., 1988; Nazzi, Bertoncini, & Mehler, 1998; Ramus et al., 2000). Linguistic rhythm may be viewed as a parameter that shows variation across the languages of the world. Three types of linguistic rhythm have been identified, leading to a classification of languages into three classes (Pike, 1945; Abercrombie, 1967; Ladefoged, 1975): stress-timed languages, including most Germanic languages as well as Russian, Arabic or Thai, syllable-timed languages, including most Romance languages as well as Turkish or Yoruba, and mora-timed languages, including Japanese. It has also been proposed that an early setting of this parameter might act as a bootstrap in the acquisition of phonology (Mehler, Dupoux, Nazzi, & Dehaene-Lambertz, 1996; Ramus, Nespor, & Mehler, 1999). However, the evidence that newborns are sensitive to linguistic rhythm does not seem to be entirely conclusive.

Some researchers have directly studied rhythm perception by infants. Demany, McKenzie, and Vurpillot (1977) showed that 2-3 month-old infants are able to discriminate sequences of tones differing in temporal organization. Fowler, Smith, and Tassinary (1986) moreover showed that 3-4 month-olds where able to discriminate sequences of syllables whose *P-centers* (Morton, Marcus, & Frankish, 1976) were isochronous or not. However, it is not clear at all whether the notion of rhythm investigated in those studies has anything to do with that of linguistic rhythm, as defined above. There are nevertheless two lines of evidence suggesting that newborns may classify languages according to their rhythm. The first one comes from the pattern of successes and failures accumulated across the different studies: newborns have been shown to discriminate between stress-timed and syllable-timed languages (French-Russian and English-Italian: Mehler et al., 1988; English-Spanish: Moon, Cooper, & Fifer, 1993) and between stress-timed and mora-timed languages (English-Japanese: Nazzi et al., 1998; Dutch-Japanese: Ramus et al., 2000). However, the only attempt to assess within-class discrimination has yielded a negative result (English-Dutch: Nazzi et al., 1998). Moreover, newborns also discriminated between a set of English and Dutch (stress-timed) sentences and a set of Spanish and Italian (syllable-timed) sentences, but failed when the two sets (English and Italian versus Dutch and Spanish) did not reflect two types of rhythm (Nazzi et al., 1998). Thus, they seem to be able to discriminate between sets of languages, if and only if these sets are congruent with different rhythm classes. As impressive as this result may be, the small num-

I thank Jacques Mehler, Marc Hauser, Anne Christophe, Christophe Pallier, Emmanuel Dupoux and Ghislaine Dehaene-Lambertz for useful discussions, Sylvie Margules and Renate Zangl for help testing the subjects, Jacques Mehler, Anne Christophe and Sarah Griffiths for comments on a previous version of this paper, Xavier Jeannin and Michel Dutat for technical assistance, and the Délégation Générale pour l'Armement for financial support.

Experiments 1 and 2 have already been partially reported in Ramus, Hauser, Miller, Morris, and Mehler (2000), to compare the perceptual abilities of newborns with those of monkeys.

Correspondence to: Franck Ramus, now at Institute of Cognitive Neuroscience, 17 Queen Square, London WC1N 3AR, GB. E-mail: f.ramus@ucl.ac.uk.

ber of languages studied does not guarantee that this pattern would hold for other unrelated languages. Indeed, considering the great variety of cues present in speech, other properties than rhythm may have allowed discrimination, and may therefore be considered as confounding factors. This concern has led researchers to look for a second line of evidence, by reducing the speech cues that were available for discrimination. Thus, Mehler et al. (1988) successfully replicated their experiments after low-pass filtering their stimuli at 400 Hz. This process, which eliminates the higher frequencies of speech, and therefore most of the phonetic information, is thought to preserve only its prosodic properties (rhythm and intonation). Similarly, the experiments by Nazzi et al. (1998) used filtered speech exclusively, and Ramus et al. (2000) used sentences that were resynthesized in such a way as to preserve only prosodic cues (see below). Thus, there is converging evidence that prosody is all newborns need to discriminate languages.

Nevertheless, prosody does not reduce to rhythm. It remains possible that its other major component, intonation, plays a role in the observed discriminations. Although we do not know of typological studies of intonation that would allow to make specific predictions for all the pairs of languages considered, it is, for instance, predictable that English and Japanese should be discriminable on the basis of their intonation. Indeed, English is a Head-Complement language, whereas Japanese is Complement-Head[1], and this syntactic parameter is said to have a prosodic correlate, *prominence*, which is signaled both in terms of rhythm and intonation (Nespor, Guasti, & Christophe, 1996). Moreover, there is empirical evidence that some languages are discriminable purely by their intonation, including English and Japanese indeed (Ramus & Mehler, 1999), English and French (Maidment, 1976, 1983) and English and Dutch (de Pijper, 1983). In order to assess whether newborns actually perceive linguistic rhythm, it is therefore necessary to get rid of the intonation confound, that is to go beyond speech filtering and remove intonation from the stimuli.

Ramus and Mehler (1999) have adapted a technique, speech resynthesis, to selectively degrade the different components of speech, including rhythm and intonation. This technique has notably been used to resynthesize different versions of English and Japanese sentences, and assess which components of speech were sufficient for discrimination of the two languages. The different versions included (a) broad phonotactics + prosody, (b) prosody, (c) rhythm only, and (d) intonation only. Results showed that pure rhythm was sufficient for French subjects to discriminate between the two languages. Pure intonation was also sufficient, but the task was more difficult and required explicit knowledge of one the two target languages. In the present series of experiments, we wish to apply the same rationale to the study of language discrimination by newborns, i.e., progressively eliminate the speech cues available for discrimination, and finally assess whether linguistic rhythm is, as hypothesized, the critical cue.

Experiment 1: Natural speech

This first experiment aims to test the discrimination of two languages in the most unconstrained condition, using natural, unsynthesized sentences. The two languages we have selected are Dutch and Japanese. The discrimination of this pair of languages was previously tested in 2-3 month-old English infants, and yielded only a marginally significant result (Christophe & Morton, 1998). This was interpreted as showing a growing focus on the native language, hence a loss of interest in foreign ones (consistent with Mehler et al., 1988). This pair of languages has never been tested on newborns, but it is expected to be easy to discriminate, given the English-Japanese discrimination by French newborns obtained by Nazzi et al. (1998), and the fact that English and Dutch are very close in many respects, including rhythm.

Materials and Method

All the experiments included in this paper use the same methodology unless otherwise stated. Since we have made special efforts to improve upon previously used procedures, our methodology is described below in great detail.

Stimuli

Dutch and Japanese sentences were taken from a corpus constituted by Nazzi (1997; Nazzi et al., 1998), comprising short news-like sentences read by four female native speakers per language[2]. We selected 5 sentences per speaker, i.e., 20 sentences per language, matched in number of syllables (15 to 19, with an average of 17) and in duration (3120 ms ± 186 for Dutch, 3040 ms ± 292 for Japanese, $F(1,39) = 1.1$, $p = 0.3$). We were also concerned about the possibility that speakers in one language might have a higher pitch than speakers in the other language. Average fundamental frequency[3] is indeed significantly different between the two languages: 216 Hz ± 19 for Dutch, 235 Hz ± 15 for Japanese, $F(1,39) = 11.8$, $p = 0.001$. This is compensated for through resynthesis in Experiment 2, and we will see that this had no influence on discrimination. Sentences in subsequent experiments were resynthesized from these 40 source sentences, and differ only with respect to the type of synthesis that was used[4].

Experimental protocol

As is customary when testing newborns, we used the non-nutritive sucking technique in a habituation paradigm (Eimas et al., 1971). Compared with previous studies (see Fernald &

---

[1] For example, relative phrases come after the corresponding verb in English, but before it in Japanese.

[2] This corpus consists exclusively of adult-directed speech.

[3] Fundamental frequency was extracted at intervals of 5 ms using the Bliss software. We calculated an average F_0 for each sentence, as the average of all its non-zero F_0 values.

[4] Samples of the different types of stimuli used in the present experiments can be heard on: http://www.ehess.fr/centres/lscp/persons/ramus/resynth/ecoute.htm.

McRoberts, 1996 for a critique), we have taken extra care to blind the experimenter with respect to the condition and to reduce direct interventions during the test to a minimum. For this purpose, the experiment has been programmed on a PC in such a way as to be maximally automatized.

Experimental conditions. Babies are randomly assigned to the control or to the experimental group. In the habituation phase, they are exposed to 10 sentences uttered by two speakers of one language. In the test phase, babies from the control group hear 10 new sentences uttered by the other two speakers in the same language, whereas babies from the experimental group hear 10 new sentences uttered by two new speakers in the other language. The language presented in the habituation phase is counterbalanced across subjects, resulting in four experimental conditions. Assignment of the subjects to the different conditions is managed by the program and withheld from the experimenter.

Procedure. The test takes place in a sound-attenuated booth, with only the experimenter and the baby inside. The experimenter sits outside of the infant's field of vision and wears a sound-proof headset playing masking noise. This noise consists of four superimposed streams of all the experimental sentences playing continuously, in order to optimally mask the stimuli. The baby is seated in a reclining chair, and is presented with a pacifier fixed on a flexible arm. The air pressure in the pacifier is measured by a pressure transducer, amplified and transmitted to the computer via an analog/digital board. The computer detects sucks and computes their relative amplitude.

During the first two minutes, the baby sucks in silence and the computer calculates a high-amplitude (HA) threshold, such that 75% of the sucks have an amplitude above the threshold. Subsequently, only HA sucks are considered[5]. The habituation phase then starts. Each HA suck may elicit one sentence, subject to the condition that only one sentence is played at once, and observing an inter stimulus interval of at least 500 ms. Sentences are played in a random order, directly from the hard disk of the computer, by two loudspeakers placed in front of the baby. The habituation phase goes on until the habituation criterion is met: it consists in a minimum 25% decrease in the number of HA sucks per minute for two consecutive minutes, compared with the maximum number of sucks previously produced in 1 minute[6]. When the criterion is met, the computer switches to the test phase, which lasts for 4 minutes. Test sentences are played in the same conditions as the habituation sentences.

Delay and rejection conditions. Other factors may interfere with the test and may make it necessary to delay the shift to the test phase or simply to discard the baby's data. We have tried as much as possible to have these decisions made automatically by the computer, on the basis of the sucking pattern and of indications given by the experimenter. When the baby loses the pacifier, starts crying, or falls asleep, the experimenter needs to take appropriate action. When doing so he presses a special key on the keyboard, indicating the occurrence of an event interfering with the baby's sucking. The most critical period in the test consists of the two minutes before and the two after the shift, which are used for the statistical analyses. It is important to ensure that during those four minutes, (a) no interference has occurred, (b) the baby was awake and sucking, and thus heard enough sentences. The computer thus implements the following conditions:

Delay:

- if some interference was signaled during the 2 minutes preceding the shift,
- OR if the baby didn't hear any sentence during any one of those 2 minutes,

then the shift is delayed for at least one minute, and the habituation phase goes on.

Rejection:

- if some interference was signaled during the 2 minutes following the shift,
- OR if the baby didn't hear any sentence during any one of those 2 minutes,
- OR if the habituation phase has already lasted for 20 minutes,

then the test is discontinued and the baby's data are discarded.

In addition, the experimenter himself may make the decision to discontinue the test, (a) if the baby refuses the pacifier, (b) if he/she doesn't keep awake or suck enough, (c) if he/she keeps crying or being agitated.

Considering that (a) the experimenter is not aware of the baby's experimental condition, (b) he can't hear the stimuli during the test, (c) decisions concerning the shift to the test phase and the acceptability of a baby's data are as automatic as possible, we regard it as an unlikely possibility that the experimenter have a significant influence (be it conscious or unconscious) on the baby's behavior, leading to biased effects.

Participants

Experiments took place at the Maternité Port-Royal, Hôpital Cochin in Paris. Participation was solicited from the mothers after birth, during their stay at the maternity hospital. One of the parents' informed consent was obtained, and experiments were run with the agreement of the CCPPRB Paris-Cochin (the hospital's ethics committee). Babies were pre-selected on the basis of their medical files, according to the following criteria:

- Age between 2 and 5 days old;
- Gestational age greater than or equal to 38 weeks;
- Birth weight greater than or equal to 2800 grams;

[5] Eliminating the weaker 25% of the sucks helps increasing the signal/noise ratio (Siqueland & DeLucia, 1969).

[6] The first minute of the phase is not taken into account for the determination of the maximum. Additional conditions impose that this maximum is at least 20 HA sucks per minute, and that the habituation phase lasts at least 5 minutes.

- No suffering at birth (APGAR score = 10 at 5 min);
- Normal medical assessments at birth and at two days;
- No seroconversion to rubella or toxoplasmosis;
- Mother not affected by viruses, and not addicted to any drug including alcohol or tobacco;
- No family history of deafness or neurological problems;
- No Dutch or Japanese spoken at home.

Babies were tested three hours after feeding on average, when they could be easily woken and kept in a quiet alert state, and when their sucking reflex was maximal.

In the present experiment, 32 babies were successfully tested, 18 males and 14 females, with a mean age of 67 ± 22 hours, a mean gestational age of $40 \pm 1;1$ weeks and a mean birth weight of 3530 ± 402 g. Twenty-nine came from monolingual French families, 2 from families where one or several other languages than French are spoken and 1 from a family where no French is spoken. The results of 42 additional babies were rejected for the following reasons: rejection of the pacifier (1), sleeping or insufficient sucking before the shift (12), crying or agitation (9), failure to meet the habituation criterion (9), sleeping or insufficient sucking after the shift (6), loss of the pacifier after the shift (4) and computer failure (1).

Results

Figure 1 shows the number of HA sucks per minute for the 2 groups of subjects. To ensure that babies were in comparable conditions during the habituation phase, an ANOVA was performed on average number of HA sucks over the 5 minutes preceding the shift, and showed no significant effect of group (control or experimental) [$F(1,31) = 2.6$, $p = 0.12$], although there is a trend for babies in the control group to suck more, and no significant effect of the language heard in habituation (Dutch or Japanese) [$F(1,31) < 1$]. In order to assess whether the experimental group reacted more to the change than the control group, we conducted an analysis of covariance (ANCOVA), with the average number of HA sucks during the 2 post-shift minutes as dependent variable, the average number of HA sucks during the 2 pre-shift minutes as covariate, and the group as independent variable[7]. Here, there is no significant group effect [$F(1,29) < 1$], showing that the babies in the experimental group have not reacted to the language change.

Discussion

This experiment shows that newborns fail to discriminate between Dutch and Japanese when the stimuli are not degraded at all, consisting just of natural sentences. This may seem inconsistent with Nazzi et al.'s (1998) finding that similar French newborns can discriminate between English and Japanese. Among the few differences between our experiment and that of Nazzi et al., is the fact that their sentences were low-pass filtered, whereas ours aren't. While it may seem that the more information, the easier the discrimination, previous experiments on adults suggest that it is not always the case: Irrelevant information may actually impair

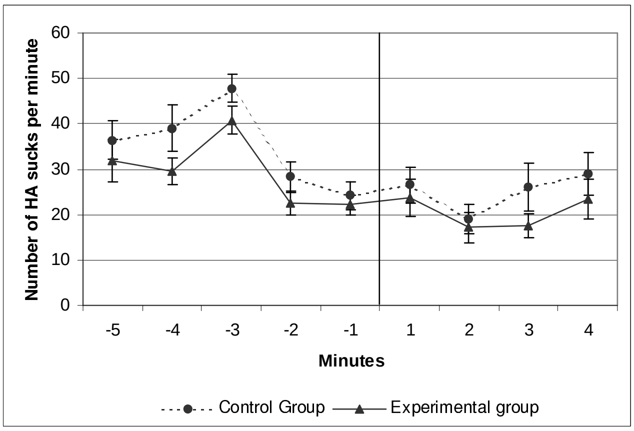

Figure 1. Exp. 1: Dutch-Japanese discrimination – Natural speech. Minutes are numbered from the shift, indicated by the vertical line. Error bars represent ± 1 standard error of the mean. Adapted with permission from Ramus et al. (2000). Copyright 2000 American Association for the Advancement of Science.

the discrimination (Ramus & Mehler, 1999; Ramus, Dupoux, Zangl, & Mehler, submitted).

Here, the fact that each newborn hears 4 speakers during the experiment may constitute such irrelevant information. Indeed, it has been suggested that the adaptation to speaker variability is a costly process in younger infants, that may interfere with other speech categorization abilities (Jusczyk, Pisoni, & Mullenix, 1992). It is actually remarkable that all language discrimination experiments performed on newborns to this day have used stimuli where speaker variability was completely absent, through the use of a single bilingual speaker (Mehler et al., 1988; Moon et al., 1993), or at least strongly attenuated through the use of low-pass filtering (Mehler et al., 1988; Nazzi et al., 1998). Experiment 1 is thus the first language discrimination experiment to expose newborns to 4 different voices.

Speech resynthesis is a convenient technique to test whether newborns were disturbed by speaker variability: whatever the original number of speakers, all sentences are synthesized using only one voice, that of the synthesizer. If our hypothesis is correct, we should then predict that newborns will be able to discriminate the two languages once the sentences are resynthesized.

[7] This is the standard analysis of sucking rates since Christophe, Dupoux, Bertoncini, and Mehler (1994) showed that it is more appropriate than doing a simple ANOVA on a dishabituation index (e.g., the difference in sucking rates between the 2 minutes after and the 2 minutes before the shift). Here, we also ran the ANOVAs and found that they always led to the same conclusions as the ANCOVAs. We thus only report the results of the latter.

Experiment 2: *Saltanaj* resynthesized speech

Materials and Method

Stimuli

We used the same sentences as for Experiment 1, but we resynthesized them in the *saltanaj* manner described by Ramus and Mehler (1999). For each sentence, the fundamental frequency (F_0) is measured, phonemes are identified and their duration measured. These parameters can be manipulated at will, and then used as input to the speech synthesizer MBROLA[8] (Dutoit, Pagel, Pierret, Bataille, & van der Vrecken, 1996). The *saltanaj* manipulation involves reducing the phonemic inventory of the sentences, by mapping each phoneme to a single instance of the same manner of articulation: fricatives are mapped to /s/, vowels to /a/, liquids to /l/, plosives to /t/, nasals to /n/ and glides to /j/. The exact phoneme durations as well the F_0 curve are copied from the original sentences. Thus, the overall rhythm and intonation of the sentences are faithfully preserved, together with broad phonotactic characteristics. However, phonetic and fine phonotactic differences are eliminated. Obviously, access to syntax and meaning is blocked as well. Voice differences are eliminated, since a single synthetic voice is used; however, prosodic differences between speakers are preserved. It therefore still makes sense to have a speaker change in the control condition, where "speakers" refers to those who produced the original sentences.

Incidentally, resynthesis gives full control over the average fundamental frequency which, we have noted earlier, was significantly different between the two languages. We have thus decided to reduce this difference by multiplying all F_0 values by 1.04 for Dutch, and by 0.96 for Japanese. Note that although this manipulation removes a possible confound, it is not a very plausible one. Indeed, in Experiment 1, the fact that Japanese speakers have a higher pitch on average isn't sufficient by itself to allow for a discrimination.

Procedure

Hesketh, Christophe, and Dehaene-Lambertz (1997), testing 2-3 month-olds with the non-nutritive sucking procedure, have improved the habituation paradigm by having each baby undergo 2 shifts (one of language, one of speakers). Thanks to the within-subject design, they found that the statistical power increased, and thus needed fewer subjects to complete an experiment.

Here we have tried to adapt this 2-shift design to experimentation on newborns. The first two phases of the experiment were run exactly as in Exp. 1. After a baby had undergone the first shift and the four-minute test phase, a second shift was made possible. It was subject to the same habituation criterion, delay and rejection conditions as the first one. When met, it allowed the baby to undergo a second four-minute test phase. In that case, the second shift was of a different kind (language or speaker) to the first one. Babies who didn't successfully pass the second shift were nevertheless kept for analysis of the first shift.

Thus, this attempt at improving the procedure did not interfere with the collection of the data on the first shift, making it possible to independently analyze the results of the two shifts. Indeed, for all babies, everything was as in Exp. 1 up until the fourth test minute. Some babies were just allowed to go on further if they could. This experiment was stopped as soon as 32 babies successfully passed the first shift.

Participants

Thirty-two babies were successfully tested, 16 males and 16 females, with a mean age of 80 ± 25 hours, a mean gestational age of $39 \pm 0;6$ weeks and a mean birth weight of 3508 ± 477 g. Twenty-five came from monolingual French families, 3 from families where one or several other languages than French are spoken and 4 from a family where no French is spoken. The results of 20 additional babies were rejected for the following reasons: sleeping or insufficient sucking before the first shift (6), crying or agitation (4), failure to meet the habituation first criterion (1), sleeping or insufficient sucking after the first shift (3), loss of the pacifier after the first shift (6).

Results

First shift

Figure 2 shows the number of HA sucks per minute for the two groups of babies. There was no significant group effect on babies' sucking during the 5 pre-shift minutes [$F(1,31) < 1$], neither was there an effect of the habituation language [$F(1,31) = 1.8$, $p = 0.19$]. An ANCOVA on the 2 post-shift minutes, controlling for the 2 pre-shift minutes, showed a significant group effect [$F(1,29) = 6.3$, $p = 0.018$], indicating that babies in the experimental group increased their sucking significantly more than those in the control group. This leads us to the conclusion that the babies in the experimental group were able to discriminate between the two languages.

Second shift

Out of 32 babies successfully passing the first shift, only 11 passed the second one. The results of the remaining 21 were discarded due to sleeping or insufficient sucking before the second shift (7), crying or agitation (2), failure to meet the second habituation criterion before the 20 minute time limit (6), sleeping or insufficient sucking after the second shift (4), loss of the pacifier after the second shift (2). The small proportion of babies able to undergo the second shift already shows that this procedure, as used on 2-month-olds, is not viable for newborns: it would lead to discard too much data (here, a total of 41 babies out of 52).

[8] MBROLA is freely available from http://tcts.fpms.ac.be/synthesis/mbrola.html.

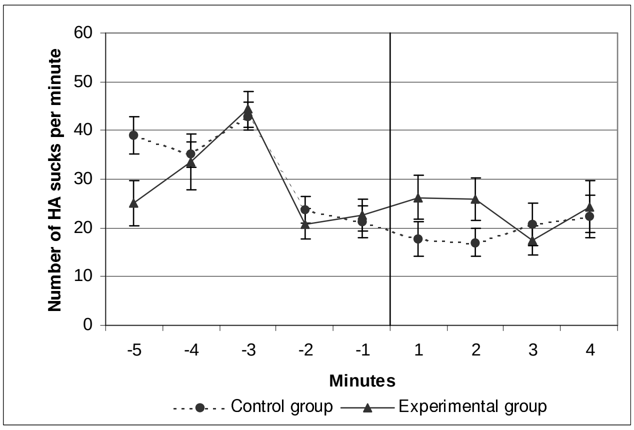

Figure 2. Exp. 2: Dutch-Japanese discrimination – *Saltanaj* speech, first shift. Minutes are numbered from the shift, indicated by the vertical line. Error bars represent ±1 standard error of the mean. Adapted with permission from Ramus et al. (2000). Copyright 2000 American Association for the Advancement of Science.

Nevertheless, a discrimination index I was computed according to the following formula (Hesketh et al., 1997):

$$I = (post2_{expe} - pre2_{expe}) - (post2_{cont} - pre2_{cont})$$

where $post2$ and $pre2$ refer to the average number of sucks during the 2 post-shift (respectively pre-shift) minutes, and where the *cont* and *expe* indices refer to the type of shift (control or experimental). Thus, a baby increasing her sucks more to the language shift than to a speaker shift will have a strictly positive I. Figure 3 gives the values of I for the 11 babies. Table 1 shows the dishabituation scores depending on type and order of the shifts.

Table 1
Dishabituation scores ($post2 - pre2$) according to order and type of the shift. The number of observations per case is indicated in parentheses.

	First shift	Second shift
control/experimental (4)	0	+0.8
experimental/control (7)	+6.6	+4.5

There is a very slight trend in the predicted direction: babies increase their sucking more during experimental shifts than during the corresponding control ones. However, this is true of 6 babies out of 11 only, and the average discrimination index is 1.9 ± 14, not significantly different from 0 [$t(10) < 1$].

Note that during the first shift, these 11 babies are behaving consistently with the others during the first shift: they increased their sucking more (+6.6) to the language change than to the speaker change (+0). Thus, whereas the first shift reveals meaningful information concerning the reactions of infants, the second shift merely shows a perseveration of the behavior produced during the first shift, which is little modulated by the nature of the second shift. In summary, this attempt suggests that there is little to learn from a second shift with newborns. However, this outcome does not diminish the results obtained on the first one, which are clear and interpretable.

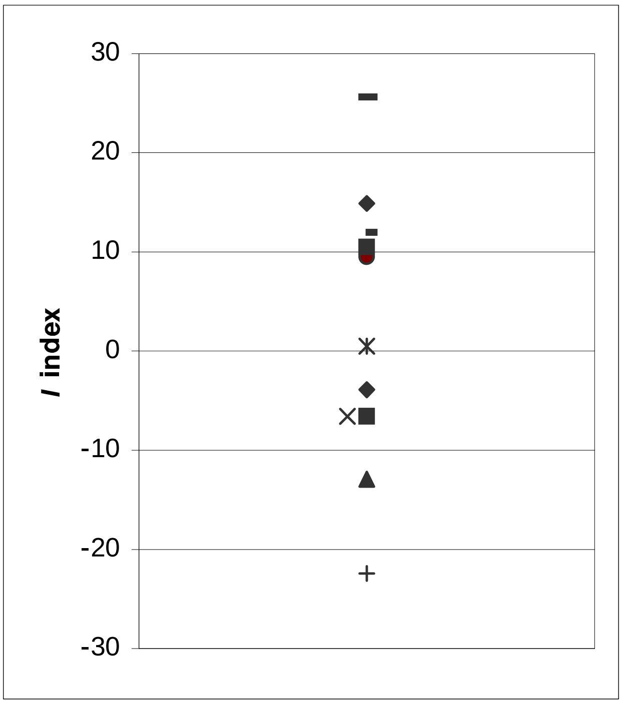

Figure 3. Exp. 2: Values of index I for 11 babies passing the two shifts.

Discussion

The data obtained on the 32 newborns who successfully passed the first shift show that (a) they are able to discriminate between Dutch and Japanese, (b) they can do so when sentences are resynthesized in the *saltanaj* manner, i.e. when lexical, syntactic, phonetic and most phonotactic information is removed..

Although the interaction with Exp. 1 is not quite significant [$F(1,59) = 2.6$, $p = 0.11$][9], this is also consistent with the hypothesis that newborns have difficulties coping with talker variability (Jusczyk et al., 1992), which would be the reason why they failed to discriminate the same sentences when they were not resynthesized.

The *saltanaj* resynthesis achieves a comparable level of stimulus degradation as low-pass filtering: Since all the durations and the fundamental frequency are faithfully reproduced, prosody, in a broad sense, is still preserved. It is

[9] Note that interaction are seldom significant in experiments on newborns anyway, due to their low statistical power. For instance, in directly comparable studies, no significant interaction were ever reported (Mehler et al., 1988; Nazzi et al., 1998).

therefore insufficient to disentangle the role of rhythm and intonation. This concern is addressed in the next two experiments.

Experiment 3: *sasasa* with artificial intonation

Materials and Method

Stimuli

In previous experiments testing language discrimination by adults on the basis of rhythm only (Ramus & Mehler, 1999; Ramus et al., submitted), sentences were resynthesized following the *flat sasasa* manner: all consonants were mapped to /s/ and all vowels to /a/, and in addition the original F_0 contour of the sentence was ignored and replaced by a constant F_0. Thus all differences concerning intonation or syllable structure were eliminated, preserving only rhythmical differences between the two languages.

When testing babies, an additional concern is to keep them awake and active in the experiment. In this respect, *flat sasasa* stimuli are potentially problematic. Both their low phonetic diversity and their monotonous intonation are susceptible to provoke boredom or distress in infants, and/or to induce them to process the stimuli as non-speech. We thus felt we had to improve the attractiveness of our stimuli, while still adequately testing our hypotheses. Considering that newborns are known to react normally to low-pass filtered speech (Mehler et al., 1988; Nazzi et al., 1998), we assumed that phonetic diversity was not a necessary condition, but we chose to preserve some variability in the intonation.

Therefore, we decided to resynthesize the same 40 sentences as before using a *sasasa* phonetic mapping, i.e., to map all consonants to /s/ and all vowels as /a/. However, instead of applying a flat intonation to each sentence, we applied artificial intonation contours. Five intonation contours inspired from French sentences were designed and each was applied to 4 of the Dutch and 4 of the Japanese sentences. All contours included a regular declination towards their end, in order to be more easily adapted to sentences of different lengths; they are illustrated in Figure 4. Thus, the resynthesized sentences incorporate both within-sentence and within-language intonational variability, but no differences in intonation between the two languages.

A potential criticism of this method is that there might be an interaction between intonation and rhythmic structure, so that the five contours selected might be more adapted to the rhythmic structure, say, of Dutch, than to that of Japanese. This would then introduce a difference between the two sets of sentences that would not be a rhythmic difference, strictly speaking. In order to investigate this possibility, we have conducted the following preliminary experiment on adult subjects.

Judgement by adult subjects of sentences resynthesized with an artificial intonation. Twelve participants were recruited and tested in a quiet room. They were 4 men and 8 women, with a mean age of 34 years, and of various native

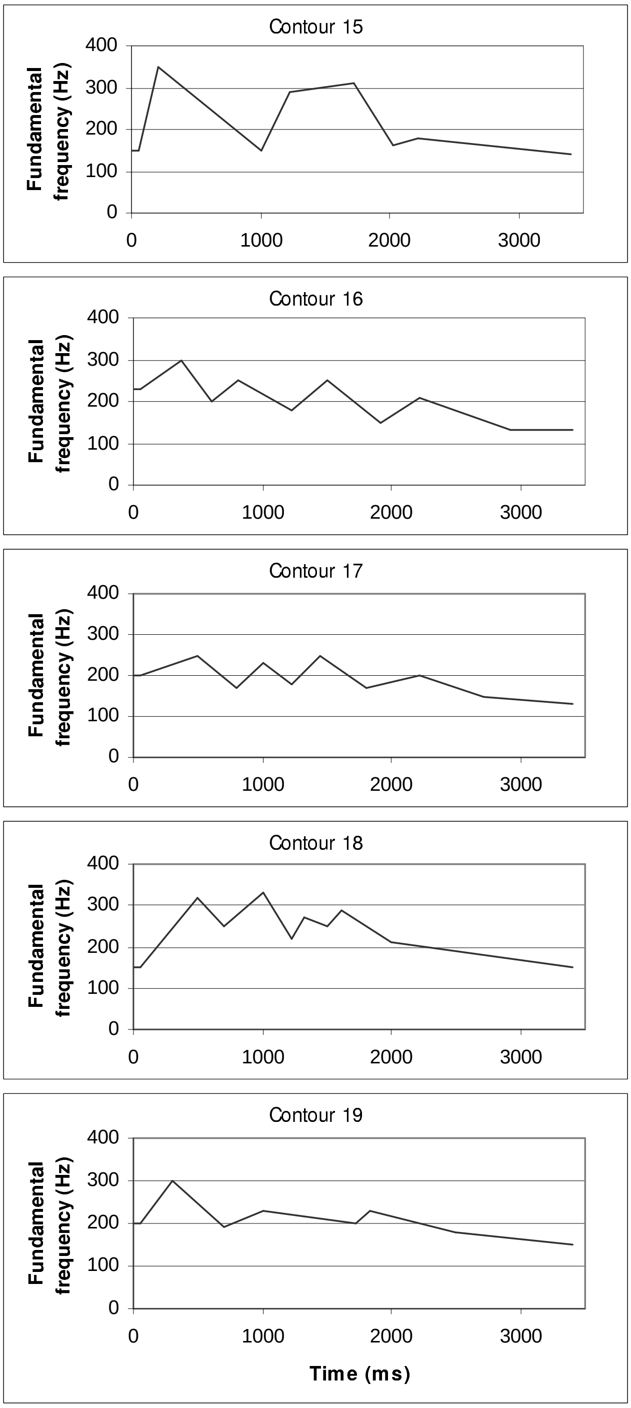

Figure 4. Intonation contours used in Exp. 3.

languages (4 French, 1 Rumanian, 3 Spanish, 1 German, 2 English, 1 Dutch).

Two blocks of stimuli were designed: one comprising the 40 experimental sentences to be used on the babies (*sasasa* with artificial intonation), and the other the *saltanaj* sentences with original intonation used in Exp. 2, to provide a baseline.

Each participant heard the sentences one by one, in a random order within each block; the order of the blocks was counterbalanced across subjects. The task was to judge how natural the intonation of each sentence was (from 0: very strange to 5: perfectly natural). If artificial intonations are equally adapted to the rhythmic structure of the two languages, they should yield similar average judgements. These are shown in Figure 5.

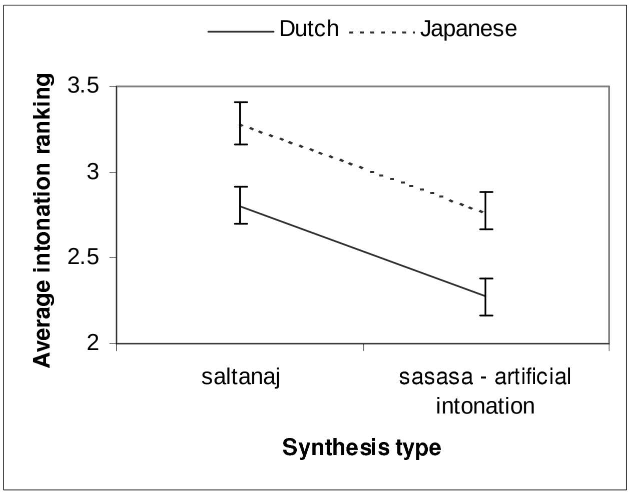

Figure 5. Adult subjects' judgements of the intonation pattern of Dutch and Japanese sentences. Error bars represent ± 1 standard error of the mean by subject.

It appears that the artificial intonations of the *sasasa* stimuli are judged to be significantly less natural in Dutch than in Japanese sentences [$F(1,11) = 12.5$, $p = 0.005$]. However, the same is true of the *saltanaj* sentences with their original intonation [$F(1,11) = 8.4$, $p = 0.02$]. It thus can't be interpreted as an effect of mismatch between the artificial intonation contours and Dutch rhythm. Rather, it seems to reflect the influence of syllabic structure on subjects' responses, although they were instructed to report specifically about intonation. From their reports, it appears that the presence of heavy consonant clusters in Dutch (also reflected by longer /s/s in the *sasasa* version) biased them in favor of Japanese. Thus, there is a main effect of language [$F(1,11) = 21.8$, $p = 0.001$], and there is also a main effect of type of synthesis [$F(1,11) = 8.4$, $p = 0.02$], revealing that *sasasa* synthesis sounded less natural to the subjects than *saltanaj*. Despite these effects that interfered with the subjects' judgements, it is most important to note that there is no interaction between language and type of stimuli [$F(1,11) < 1$], indicating that the artificiality of the *sasasa* stimuli's intonation did not interact with the rhythmic structure of the two languages[10], which was the hypothesis under test. We therefore consider our stimuli as appropriate to test language discrimination on the basis of rhythm only.

Procedure

Due to the unsuccessful attempt to have the babies undergo 2 shifts in Exp. 2, we abandoned the second shift and reverted to the one-shift procedure used in Exp. 1.

Participants

Forty babies[11] were successfully tested, 19 males and 21 females, with a mean age of 66 ± 23 hours, a mean gestational age of $40;1 \pm 1;1$ weeks and a mean birth weight of 3428 ± 424 g. Twenty-six came from monolingual French families, 11 from families where one or several other languages than French are spoken and 3 from families where no French is spoken. The results of 20 additional babies were rejected for the following reasons: rejection of the pacifier (2), sleeping or insufficient sucking before the shift (7), crying or agitation (3), failure to meet the habituation criterion (1), sleeping or insufficient sucking after the shift (4), loss of the pacifier after the shift (3).

Results

Figure 6 shows the number of HA sucks per minute for the two groups of babies. There was no significant group effect on babies' sucking during the 5 pre-shift minutes [$F(1,39) = 1.5$, $p = 0.23$]. However, there was a significant effect of the habituation language [$F(1,39) = 12.8$, $p = 0.001$], babies listening to Dutch sucking more (35.6 ± 8.3 sucks per minute on average) than those listening to Japanese (25.8 ± 9.1 sucks per minute). To take this effect into account, we included the habituation language factor in the usual ANCOVA. We found that it had no significant effect on sucking patterns during the 2 post-shift minutes [$F(1,35) < 1$], and most importantly, that it did not interact with the group factor [$F(1,35) < 1$]. Thus, this effect had no consequence on the overall result of the experiment; its interpretation will be addressed in a later section. Finally, the ANCOVA on the 2 post-shift minutes, controlling for the 2 pre-shift minutes, showed no group effect [$F(1,35) < 1$], indicating that the babies didn't discriminate between the two languages.

Discussion

There are several possible interpretations of the failure of babies to discriminate between Dutch and Japanese given the *sasasa* sentences with artificial intonation. One is that babies don't process speech rhythm, and that language discrimination experiments should be re-interpreted as revealing the processing of intonation. Another is that rhythm and

[10] Note also that no floor nor ceiling effect would have prevented this interaction to appear.

[11] Eight additional babies were tested after analyses on the first 32 babies were found a little ambiguous.

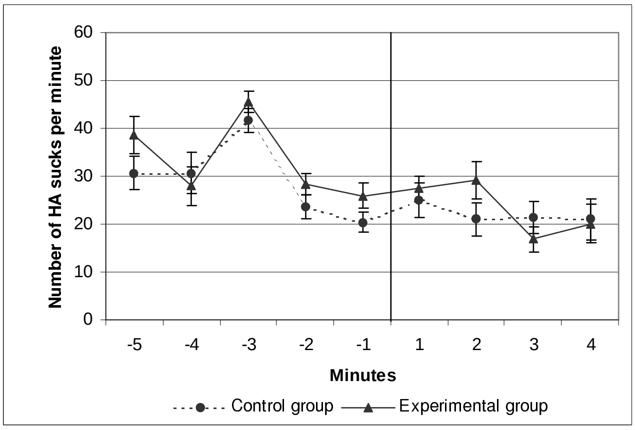

Figure 6. Exp. 3: Dutch-Japanese discrimination – *Sasasa* speech with artificial intonation. Minutes are numbered from the shift, indicated by the vertical line. Error bars represent ±1 standard error of the mean.

intonation are inseparable: babies may be sensitive to speech rhythm, but intonation is necessary to fully process it. Yet another interpretation would be that babies can process speech rhythm, but the stimuli used were inadequate for them to exhibit this ability.

For instance, Ramus et al. (2000) showed that newborns don't discriminate Dutch from Japanese anymore when the same sentences are played backwards. This suggests that stimuli that are not enough speech-like are not correctly processed by babies, even though they contain enough basic acoustic information for the discrimination to be feasible in principle. In this respect, *sasasa* might not be speech-like enough: there is indeed no natural language with so little phonetic diversity. It could also be that these stimuli are too boring or distressing for the babies, as we had hypothesized regarding a flat intonation. Whatever the appropriate explanation, we will now try to increase the chances that the babies correctly process the stimuli.

Experiment 4: *saltanaj* with artificial intonation

Materials and Method

Stimuli

There are two differences between the stimuli used in Experiment 2 and those of Experiment 3: one is the reduction of the phonetic inventory (from *saltanaj* to *sasasa*), and the other is the use of artificial intonation contours instead of the original ones. At least one of them has caused babies to fail in the discrimination task. It is therefore natural to undo one of those changes in order to know which was critical. We thus reverted to the *saltanaj* stimuli of Exp. 2, but this time we applied them the artificial intonation contours of Exp. 3.

Participants

Forty babies were successfully tested, 21 males and 19 females, with a mean age of 68 ± 21 hours, a mean gestational age of $40;1 \pm 1$ weeks and a mean birth weight of 3512 ± 341 g. Twenty-seven came from monolingual French families, 12 from families where one or several other languages than French are spoken and one from a family where no French is spoken. The results of 44 additional babies were rejected for the following reasons: rejection of the pacifier (5), sleeping or insufficient sucking before the shift (22), crying or agitation (8), failure to meet the habituation criterion (2), sleeping or insufficient sucking after the shift (4), loss of the pacifier after the shift (3).

Results

Figure 7 shows the number of HA sucks per minute for the two groups of babies. There was no significant group effect on babies' sucking during the 5 pre-shift minutes [$F(1,39) < 1$], neither was there an effect of the habituation language [$F(1,39) = 2.1$, $p = 0.16$]. An ANCOVA on the 2 post-shift minutes, controlling for the 2 pre-shift minutes, shows no significant group effect [$F(1,37) = 1.46$, $p = 0.24$]. However, examination of Figure 7 suggests that there is an effect, which is confined to the first minute after the shift. A new ANCOVA, taking as dependent variable the number of sucks during the first post-shift minute, and controlling for the 2 pre-shift minutes, yields a significant group effect indeed [$F(1,37) = 4.48$, $p = 0.04$]. This suggests that the newborns have again discriminated between Dutch and Japanese. However, the effect is weaker than in Experiment 2, being evident during only one minute following the language change.

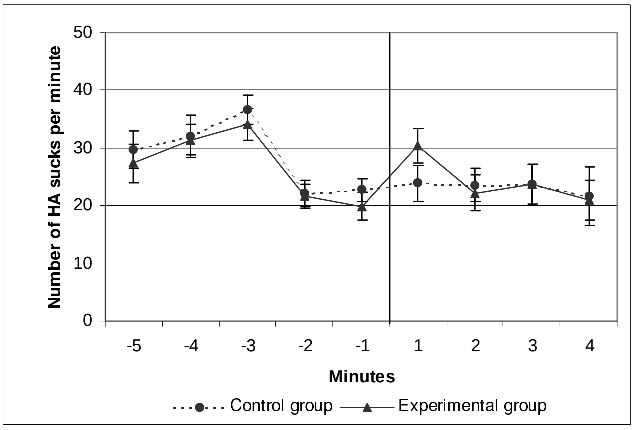

Figure 7. Exp. 4: Dutch-Japanese discrimination – *Saltanaj* speech with artificial intonation. Minutes are numbered from the shift, indicated by the vertical line. Error bars represent ±1 standard error of the mean.

Discussion

We can now come back to the three different interpretations of Experiment 3 we have proposed. Since no intonational difference remained between the two languages in the present *saltanaj* stimuli, intonation is not likely to be the cue whose processing is responsible for the language discrimination patterns observed. For the same reason, intonation cannot be necessary for rhythm processing. The most plausible interpretation is therefore that newborns can process speech rhythm, but something in the *sasasa* stimuli prevented them from correctly processing it. For instance, their extremely low phonetic diversity might lead babies to process them as non-speech. Alternatively, the constant frication of *sasasa* sentences may have been too distressing for babies to correctly perform the task; indeed, adult subjects also rated these stimuli lower than the *saltanaj* in Exp. 3, and complained about the harsh sound of the /s/s. Whatever the correct explanation may be, the alternatives could be tested by running further discrimination experiments while manipulating the nature and the variety of the phonemes used in the resynthesis.

While it is clear now that intonation does not play a crucial role in newborns' ability to discriminate between languages, the last two experiments are still open to an alternative interpretation: that newborns discriminate the *saltanaj* stimuli on the basis of basic phonotactic differences between Dutch and Japanese (e.g., the presence/absence of consonant clusters); since these differences are not preserved in the *sasasa* version, this would explain both failure in Exp. 3 and success in Exp. 4. There are, however, good independent reasons to doubt that newborns are sensitive to these phonotactic differences. Indeed, this has been directly tested in experiments where newborns had to discriminate between lists of words of different syllabic structure. For instance, newborns were unable to discriminate bi-syllabic words with complex syllabic structure (e.g., CVCCCV, CCVCCV, CVCCVC...) from bi-syllabic words with simple syllabic structure (e.g., CVCV, VCCV, VCVC...), although they were able to discriminate simple bi-syllabic (CVCV) from tri-syllabic words (CVCVCV) (Bijeljac-Babic et al., 1993). Similarly, they were unable to discriminate between tri-moraic (CVCCV) and bi-moraic words (CVCV), although they were able to discriminate again bi-syllabic from tri-syllabic words (Bertoncini et al., 1995). If newborns were able to extract phonotactic regularities from 3-second long Dutch and Japanese utterances, they would be expected to do so also from bi-syllabic words. The fact that they do not suggests that sensitivity to phonotactic differences is not available at birth; this is consistent with evidence that familiarity with the native language's phonotactic pattern emerges between 6 and 9 months of age (Jusczyk, Friederici, Wessels, Svenkerud, & Jusczyk, 1993; Friederici & Wessels, 1993; Jusczyk, Luce, & Charles-Luce, 1994). In summary, the most plausible interpretation of newborn's ability to discriminate Dutch from Japanese in their *saltanaj* version with artificial intonation, is that they have reacted to the rhythmic differences between the two languages.

Post-hoc analysis: Preference for Dutch

After finding a significant effect of language during the habituation phase of Exp. 3 (with babies sucking more to Dutch than to Japanese), we have looked for a similar trend in the other experiments. As shown in Table 2 and Figure 8, the same trend was present in all four experiments, suggesting a consistent phenomenon.

Table 2
Average number of HA sucks produced during the 5 pre-shift minutes, as a function of Experiment and language.

	Exp. 1	Exp. 2	Exp. 3	Exp. 4
Dutch	33 ± 11	32.2 ± 11.2	35.6 ± 8.3	29.4 ± 9.6
Japanese	30.9 ± 11.1	27 ± 10.4	25.8 ± 9.1	25.3 ± 8.3

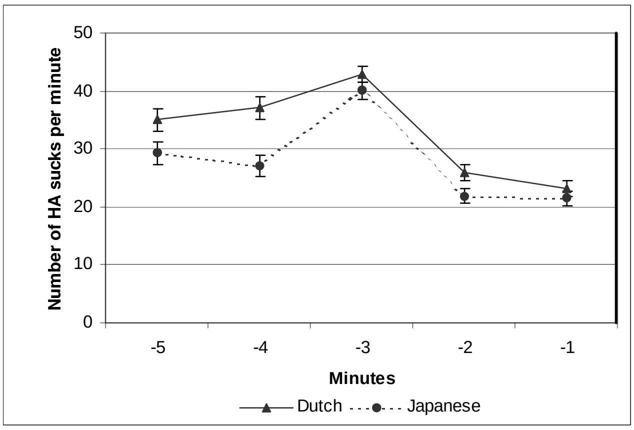

Figure 8. Average number of HA sucks produced during the 5 pre-shift minutes, as a function of language. Error bars represent ± 1 standard error of the mean.

Overall, during the 5 minutes preceding the shift, babies produced on average 32.5 ± 10 HA sucks per minute to listen to Dutch, and 27.1 ± 9.7 to listen to Japanese. The difference is significant: $F(1,142) = 11.1$, $p = 0.001$. The fact that this effect is consistent across 4 experiments suggest that this is not a random sampling effect. It is remarkable that a similar effect was also noted by Nazzi et al. (1998): they found that French newborns sucked more to listen to English than to listen to Japanese. In earlier studies, such patterns had been interpreted as showing babies' "preference" for their native-language: Mehler et al. (1988) found that French newborns sucked more during the habituation phase to listen to French than to Russian, and Moon et al. (1993), using a technique directly assessing preference, found that newborns sucked more to listen to their native language (English or Spanish). Similar results have been found in older babies as well (Dehaene-Lambertz & Houston, 1998; Bosch & Sebastián-Gallés, 1997).

Here (and also in Nazzi et al.'s study), where neither language was the babies' native language, one may want to look for an alternative explanation. For instance, following adult subjects' judgements in Exp. 3, one might argue that Dutch sentences, containing more consonant clusters and more frication, are less pleasant to listen to, and thus keep the babies more excited. Such an interpretation would, however, assume that stimuli that are less preferred provoke more sucking, which would appear to be in contradiction with the earlier studies' interpretations. In the absence of a good model of what provokes a baby's sucks, the question remains open[12].

It is yet possible to provide an interpretation of the present results in terms of genuine preference. Indeed, French can be seen as closer to Dutch and English than to Japanese along a number of dimensions. Regarding the most relevant one, rhythm, objective acoustic/phonetic measures of rhythmic properties suggest that French rhythm is much closer to that of Dutch and English than to that of Japanese (Ramus et al., 1999). Similar arguments could be made for syllable structure, intonation, size of the phonemic repertoire... It is thus conceivable that Dutch and English sound more familiar to the French newborn than Japanese. Native-language preference might therefore be re-interpreted as preference for the most familiar stimulus (along the dimensions that are relevant to the baby).

Conclusion

The literature on infant speech perception has suggested for years that newborns discriminate languages' rhythmic patterns. The evidence accumulated so far, although compelling, has always left open the possibility that the discriminations observed might be due to a sensitivity to intonational differences. Here, using resynthesized stimuli, we have shown that intonation is not necessary for newborns to discriminate between Dutch and Japanese. This does not strictly imply that intonation was not used by babies to discriminate languages in previous experiments, but it makes the "intonation hypothesis" more unlikely than ever, thereby strengthening the "rhythm hypothesis".

The point of the present study is not, however, to deny that intonation can be processed by newborns and play a role in language acquisition (see for instance Guasti, Nespor, Christophe, & van Ooyen, in press), but to show that rhythm is sufficient for babies to discriminate languages, and therefore that they genuinely process rhythm. This in turn reinforces the plausibility of language acquisition scenarios that rely on rhythm as a cue to other properties of language the infant has to learn. Noticing the correlation between a language's rhythm type and the structure of its syllables, Ramus et al. (1999) proposed for instance a detailed scenario of how the early perception of rhythm might cue the acquisition of syllable structure. It has also been proposed that procedures that listeners use to normalize for speech rate are language-specific, and indeed depend on the type of rhythm of a given language (Pallier, Sebastian-Galles, Dupoux, Christophe, & Mehler, 1998; Sebastián-Gallés, Dupoux, & Costa, in press). Since speech rate variability, like talker variability (see Exp. 1), is likely to be an obstacle to the extraction of relevant linguistic units by the infant, an interesting hypothesis would be that the perception of speech rhythm allows for an early selection of the appropriate normalization strategy. The "syllable structure" and the "normalization" hypotheses are obviously independent and not exclusive. More generally, it is conceivable that speech rhythm may allow infants to make quite a large array of language-specific adjustments on their phonological representations.

"Phonological bootstrapping" scenarios (Morgan & Demuth, 1996) like those proposed above typically require three lines of evidence: (a) a correlation between a perceptible acoustic/phonological cues and a linguistic property to be learnt, (b) evidence that the acoustic/phonological cues are processed by infants sufficiently early, and (c) evidence that infants actually use those cues in their acquisition of the linguistic property. The first line of evidence typically comes from linguistic and experimental studies like those of Ramus et al. (1999), Pallier et al. (1998) and Sebastián-Gallés et al. (in press). The second line comes from perceptual studies conducted on infants like the present one. Here, we have shown that speech rhythm is indeed a cue that is processed by newborns. The third line of evidence may be more difficult to obtain, because the actual causes of observed developmental changes are never certain. The proposed scenarios yet predict that an impairment in speech rhythm perception may delay or impair the infant's acquisition of correct syllable structure and/or speech rate normalization.

Such predictions may eventually be testable if suitable neuropsychological cases are found. We speculate that dyslexic children might provide us with such a case. Although speech rhythm perception has never been specifically tested in dyslexics, there is some evidence that they may have trouble discriminating the rhythmic patterns of sequences of tones (McGivern, Berka, Languis, & Chapman, 1991; Kujala et al., 2000). Similarly, neither the representation of syllable structure nor speech rate normalization have been investigated in dyslexics. However, informal observations point at particular difficulties with handling complex consonant clusters, hence suggesting a possibly impoverished representation of syllables. Dyslexics, or alternatively other language-impaired populations, might thus provide a critical test of the hypothesized bootstrapping role of rhythm.

References

Abercrombie, D. (1967). *Elements of general phonetics.* Chicago: Aldine.

Bertoncini, J., Floccia, C., Nazzi, T., & Mehler, J. (1995). Morae and syllables: Rhythmical basis of speech representations in neonates. *Language and Speech, 38,* 311-329.

Bertoncini, J., & Mehler, J. (1981). Syllables as units in infant perception. *Infant Behavior and Development, 4,* 247-260.

[12] Note that this problem is not particular to the sucking behavior. In studies using preferential listening techniques, preference sometimes goes for the novel stimulus, and sometimes for the familiar one. No generalized account of infants' preferences has been proposed.

Bijeljac-Babic, R., Bertoncini, J., & Mehler, J. (1993). How do four-day-old infants categorize multisyllabic utterances? *Developmental Psychology*, *29*, 711-721.

Bosch, L., & Sebastián-Gallés, N. (1997). Native language recognition abilities in 4-month-old infants from monolingual and bilingual environments. *Cognition*, *65*, 33-69.

Christophe, A., Dupoux, E., Bertoncini, J., & Mehler, J. (1994). Do infants perceive word boundaries? An empirical study of the bootstrapping of lexical acquisition. *Journal of the Acoustical Society of America*, *95*(3), 1570-1580.

Christophe, A., & Morton, J. (1998). Is Dutch native English? Linguistic analysis by 2-month-olds. *Developmental Science*, *1*(2), 215-219.

de Pijper, J. R. (1983). *Modelling British English intonation.* Dordrecht - Holland: Foris.

Dehaene-Lambertz, G., & Houston, D. (1998). Faster orientation latencies toward native language in two-month old infants. *Language and Speech*, *41*(1), 21-43.

Demany, L., McKenzie, B., & Vurpillot, E. (1977). Rhythm perception in early infancy. *Nature*, *266*(5604), 718-9.

Dutoit, T., Pagel, V., Pierret, N., Bataille, F., & van der Vrecken, O. (1996). The MBROLA Project: Towards a set of high-quality speech synthesizers free of use for non-commercial purposes. In *ICSLP'96.* Philadelphia.

Eimas, P., Siqueland, E., Jusczyk, P., & Vigorito, J. (1971). Speech perception in infants. *Science*, *171*, 303-306.

Fernald, A., & McRoberts, G. (1996). Prosodic bootstrapping: A critical analysis of the argument and the evidence. In J. L. Morgan & K. Demuth (Eds.), *Signal to Syntax: Bootstrapping from Speech to Grammar in Early Acquisition* (p. 365-388). Mahwah, NJ: Lawrence Erlbaum Associates.

Fowler, C. A., Smith, M. R., & Tassinary, L. G. (1986). Perception of syllable timing by prebabbling infants. *Journal of the Acoustical Society of America*, *79*(3), 814-825.

Friederici, A. D., & Wessels, J. M. I. (1993). Phonotactic knowledge of word boundaries and its use in infant speech perception. *Perception & Psychophysics*, *54*(3), 287-295.

Guasti, M. T., Nespor, M., Christophe, A., & van Ooyen, B. (in press). Pre-lexical setting of the head-complement parameter through prosody. In J. Weissenborn & B. Höhle (Eds.), *How to get into language: Approaches to bootstrapping in early language development.* Amsterdam: Benjamins.

Hesketh, S., Christophe, A., & Dehaene-Lambertz, G. (1997). Nonnutritive sucking and sentence processing. *Infant Behavior and Development*, *20*, 263-269.

Jusczyk, P., Luce, P., & Charles-Luce, J. (1994). Infants' sensitivity to phonotactic patterns in the native language. *Journal of Memory and Language*, *33*, 630-645.

Jusczyk, P. W., Friederici, A., Wessels, J., Svenkerud, V., & Jusczyk, A. (1993). Infants' sensitivity to the sound pattern of native language words. *Journal of Memory and Language*, *32*, 402-420.

Jusczyk, P. W., Pisoni, D. B., & Mullenix, J. (1992). Some consequences of stimulus variability on speech processing by 2-month-old infants. *Cognition*, *43*, 253-291.

Kujala, T., Myllyviita, K., Tervaniemi, M., Alho, K., Kallio, J., & Näätänen, R. (2000). Basic auditory dysfunction in dyslexia as demonstrated by brain-activity measurements. *Psychophysiology*, *37*, 262-266.

Ladefoged, P. (1975). *A course in phonetics.* New York: Harcourt Brace Jovanovich.

Maidment, J. A. (1976). Voice fundamental frequency characteristics as language differentiators. *Speech and hearing: Work in progress, University College London*, 74-93.

Maidment, J. A. (1983). Language recognition and prosody: further evidence. *Speech, hearing and language: Work in progress, University College London*, *1*, 133-141.

McGivern, R., Berka, C., Languis, M., & Chapman, S. (1991). Detection of deficits in temporal pattern discrimination using the seashore rhythm test in young children with reading impairments. *J. Learn. Disabil.*, *24*(1), 58-62.

Mehler, J., Dupoux, E., Nazzi, T., & Dehaene-Lambertz, G. (1996). Coping with linguistic diversity: The infant's viewpoint. In J. L. Morgan & K. Demuth (Eds.), *Signal to Syntax: Bootstrapping from Speech to Grammar in Early Acquisition* (p. 101-116). Mahwah, NJ: Lawrence Erlbaum Associates.

Mehler, J., Jusczyk, P., Lambertz, G., Halsted, N., Bertoncini, J., & Amiel-Tison, C. (1988). A precursor of language acquisition in young infants. *Cognition*, *29*, 143-178.

Moon, C., Cooper, R. P., & Fifer, W. P. (1993). Two-day-olds prefer their native language. *Infant Behavior and Development*, *16*, 495-500.

Morgan, J., & Demuth, K. (1996). Signal to syntax: An overview. In J. Morgan & K. Demuth (Eds.), *Signal to Syntax: Bootstrapping from speech to grammar in early acquisition* (p. 1-22). Mahwah, NJ: Lawrence Erlbaum Associates.

Morton, J., Marcus, S., & Frankish, C. (1976). Perceptual centers (P-centers). *Psychological Review*, *83*(5), 405-408.

Nazzi, T. (1997). *Du rythme dans l'acquisition et le traitement de la parole.* Unpublished doctoral dissertation, Ecole des Hautes Etudes en Sciences Sociales.

Nazzi, T., Bertoncini, J., & Mehler, J. (1998). Language discrimination by newborns: towards an understanding of the role of rhythm. *Journal of Experimental Psychology: Human Perception and Performance*, *24*(3), 756-766.

Nespor, M., Guasti, M. T., & Christophe, A. (1996). Selecting word order: The rhythmic activation principle. In U. Kleinhenz (Ed.), *Interfaces in Phonology* (p. 1-26). Berlin: Akademie Verlag.

Pallier, C., Sebastian-Galles, N., Dupoux, E., Christophe, A., & Mehler, J. (1998). Perceptual adjustment to time-compressed speech: a cross-linguistic study. *Memory & Cognition*, *26*, 844-851.

Pike, K. L. (1945). *The intonation of American English.* Ann Arbor, Michigan: University of Michigan Press.

Ramus, F., Dupoux, E., Zangl, R., & Mehler, J. (submitted). An empirical study of the perception of language rhythm.

Ramus, F., Hauser, M. D., Miller, C., Morris, D., & Mehler, J. (2000). Language discrimination by human newborns and by cotton-top tamarin monkeys. *Science*, *288*(5464), 349-351.

Ramus, F., & Mehler, J. (1999). Language identification with suprasegmental cues: A study based on speech resynthesis. *Journal of the Acoustical Society of America*, *105*(1), 512-521.

Ramus, F., Nespor, M., & Mehler, J. (1999). Correlates of linguistic rhythm in the speech signal. *Cognition*, *73*(3), 265-292.

Sebastián-Gallés, N., Dupoux, E., & Costa, A. (in press). Adaptation to time-compressed speech: Phonological determinants. *Perception & Psychophysics.*

Siqueland, E., & DeLucia, C. (1969). Visual reinforcement of nonnutritive sucking in human infants. *Science*, *165*(898), 1144-6.

van Ooyen, B., Bertoncini, J., Sansavini, A., & Mehler, J. (1997). Do weak syllables count for newborns? *Journal of the Acoustical Society of America*, *102*(6), 3735-3741.