

8-31-2020

A Notion of Prominence for Games with Natural-Language Labels

Alessandro Sontuoso

Chapman University, sontuoso@chapman.edu

Sudeep Bhatia

University of Pennsylvania, bhatisu@sas.upenn.eduFollow this and additional works at: https://digitalcommons.chapman.edu/esi_working_papersPart of the [Econometrics Commons](#), [Economic Theory Commons](#), and the [Other Economics Commons](#)

Recommended Citation

Sontuoso, A., & Bhatia, S. (2020). Notion of prominence for games with natural-language labels. *ESI Working Paper 20-38*. https://digitalcommons.chapman.edu/esi_working_papers/333/

This Article is brought to you for free and open access by the Economic Science Institute at Chapman University Digital Commons. It has been accepted for inclusion in ESI Working Papers by an authorized administrator of Chapman University Digital Commons. For more information, please contact laughtin@chapman.edu.

A Notion of Prominence for Games with Natural-Language Labels

Comments

ESI Working Paper 20-38

A Notion of Prominence for Games with Natural-Language Labels

Alessandro Sontuoso *and* Sudeep Bhatia

August 31, 2020

Abstract. We study games with natural-language labels (i.e., strategic problems where options are denoted by words), for which we propose and test a measurable characterization of prominence. We assume that – *ceteris paribus* – players find particularly prominent those strategies that are denoted by words more frequently used in their everyday language. To operationalize this assumption, we suggest that the prominence of a strategy-label is correlated with its frequency of occurrence in large text corpora, such as the Google Books corpus (“n-gram” frequency). In testing for the strategic use of word frequency, we consider experimental games with different incentive structures (such as incentives *to* and *not to* coordinate), as well as subjects from different cultural/linguistic backgrounds. Our data show that frequently-mentioned labels are more (less) likely to be selected when there are incentives to match (mismatch) others. Furthermore, varying one’s knowledge of the others’ country of residence significantly affects one’s reliance on word frequency. Overall, the data show that individuals play strategies that fulfill our characterization of prominence in a (boundedly) rational manner.

KEYWORDS: focal points; salience; coordination; hide-and-seek; culture; language; level-*k*.

JEL Classification Numbers: C72; C91.

Contact: Alessandro Sontuoso, *Smith Institute for Political Economy and Philosophy*, Chapman University, 303 W. Palm Ave., Orange, California, 92866 and *Philosophy, Politics and Economics*, University of Pennsylvania, 249 S. 36th St., Philadelphia, Pennsylvania, 19104, sontuoso@sas.upenn.edu; Sudeep Bhatia, *Department of Psychology*, University of Pennsylvania, 425 S. University Ave., Philadelphia, Pennsylvania, 19104, bhatiasu@sas.upenn.edu.

I - Introduction

Coordination problems can affect a broad array of economic interactions, yet they are often solved by shared *cultural* understandings. Culture is in fact believed to facilitate coordination by informally codifying the common beliefs and practices that are shared by a society's members (Alesina and Giuliano, 2015). For example, global companies that wish to match potential demand in developing markets usually adjust their strategies in such a way to connect to (and hence coordinate with) culturally diverse consumers. In that case, cultural awareness affects an organization's marketing activities – from labeling to customer services – thereby making *some* products more salient for *some* consumer segments (Kapferer, 2012). More generally, it has been suggested that cultural and sociolinguistic competence favor the coordination of economic interactions by promoting trade, development, and growth (Easterly and Levine, 1997; Lazear, 1999; Melitz, 2008). Natural languages are indeed an instance of a cultural code that eases the exchange of information and hence facilitates economic activities (Ginsburgh and Weber, 2020).¹

Thomas Schelling (1960) was the first to note that a wide range of economic interactions can be formally represented as a “coordination problem”, that is, a symmetric simultaneous-move game with multiple pure-strategy Nash equilibria.² Schelling informally observed that coordination puzzles are often solved by exploiting *contextual* or *cultural* cues, that drive expectations in such a way to make a specific course of action salient in specific circumstances. As we will elaborate below, more recent explanations of prominence (i.e., salience) have centered around features of the game that players perceive as distinctive, where such distinctive

¹ Research in political economy has shown that countries' per-capita GDP growth is inversely related to their ethno-linguistic fractionalization (Alesina, Devleeschauwer, Easterly, Kurlat, and Wacziarg, 2003). Also, it has been suggested that production line workers coordinate tasks more efficiently if they speak the same language, although this might result in a form of labor market discrimination against cultural minorities (Lang, 1986; Kossoudji, 1988).

² The class of coordination problems contains any situation in which there are multiple ways agents may “match” their behavior for mutual benefit. This class contains a vast and diverse array of interactions, including games with (slightly) conflicting interests, with and without Pareto-rankable equilibria. Specifically, interactions where players wish to coordinate – but have conflicting interests – are sometimes referred to as *impure* coordination games (e.g., Luce and Raiffa's, 1957, battle of the sexes is a classic example). Bacharach's (2006) Hi-Lo game is an instance of a conflict-free problem with two equilibria, one of which is Pareto-dominated by the other. On the other hand, Schelling's (1960) driving game is a case where each player is completely indifferent between the equilibria. See Lewis (1969) for an early book-length account of coordination games.

features do not necessarily vary across cultures.³ Here, on the other hand, we propose and test a characterization of prominence that directly rests on players' culture.

To that end, we note that advances in online data collection and processing have made it possible for researchers to access (culture-dependent) digitized information at little cost. These advances have boosted the use of “big data” to investigate economic, psychological, and medical outcomes (e.g., Varian, 2014; Einav and Levin, 2014; Matz and Netzer, 2017; Yin, Sulieman, and Malin, 2019). In particular, recent developments have facilitated the extraction of information from online text documents, which has made it possible to track language use across different cultures and over long periods of time (Michel et al., 2011). Such developments have led us to propose the following notion of prominence for “games with natural-language labels” (i.e., strategic problems where options are denoted by words).

In a nutshell, we assume that *players find particularly prominent those strategies that are denoted by words more frequently used in the players' everyday language, ceteris paribus*. Specifically, to operationalize this assumption we suggest that the prominence of a strategy-label is correlated with its frequency of occurrence in large text corpora, like the Google Books corpus (“n-gram” frequency).⁴ In testing for the strategic use of word frequency, we consider experimental games with different incentive structures (such as incentives *to* and *not to* coordinate), participants from different cultural backgrounds, as well as participants with asymmetric knowledge about the counterpart's cultural background.

Before fleshing out our hypotheses and methods, we note that formal (game-theoretic) characterizations of prominence have typically revolved around features of the strategic problem, that would drive one's perception as to the uniqueness of a solution independently of one's culture. For example, Harsanyi and Selten's (1988) payoff-dominance criterion assumes that *any* rational player (upon facing a one-shot game with no pre-play communication) would discard solutions that are collectively suboptimal: in that case, the payoff structure of the game serves as a “cue”, thereby directing players' expectations toward the collectively optimal solution. Related characterizations of prominence have integrated the payoff-dominance criterion into a theory of

³ For example, Bacharach and Bernasconi (1997) noted that players solve coordination games by exploiting distinctive attributes of the strategy space. The authors went on to note that – in the class of games where strategy options are represented by visual objects – there «are attributes, such as colors, comparative sizes, and simple geometric shapes, whose saliencies are *universal* constants [...]. But for others, saliencies are *culture-dependent*» (Bacharach and Bernasconi, 1997, p. 7, our italics).

⁴ This is a standard metric for word frequency in languages, and has so far been used for psychological, sociological, and historical research (e.g., Hills, Proto, Sgroi, and Seresinhe, 2019; Garg, Schiebinger, Jurafsky, and Zou, 2018).

framing, whereby players select the collectively optimal solution to some perceptual description of the game.⁵ In this respect, lab experiments have provided evidence confirming that *payoff* and *frame* asymmetries do – each to a different extent – affect behavior. That is, experiments have shown that if one of the game solutions is an “oddy” (in the sense that its label or payoff profile differ from the others), then the distinctiveness of that solution serves as a cue to facilitate coordination (e.g., Bacharach and Bernasconi, 1997; Rubinstein, Tversky, and Heller, 1997; Crawford, Gneezy, and Rottenstreich, 2008). In summary, per the above explanations of prominence, a solution is generally viewed as a “focal point” by virtue of culture-invariant cues. In what follows, instead, we explicitly study culture’s role in affecting behavior by introducing an a-priori measurable proxy for prominence (i.e., word frequency).

Our approach rests on a psychologically grounded characterization of focality, that has important implications for our understanding of strategic reasoning. Research in cognitive psychology has shown that the *frequency of exposure* to words is closely related to word *fluency*, that is, the ease with which an individual is able to recognize, retrieve, and process a word. Word frequency – through its effect on fluency – has been shown to have a role in a wide range of memory and language tasks (e.g., Anderson and Schooler, 1991; Balota and Spieler, 1999; Jescheniak and Levelt, 1994; Morrison and Ellis, 1995; Seidenberg and McClelland, 1989). Notably, there is evidence that individuals use word frequency as a cue in several *non*-strategic domains, like probability judgment (e.g., Dougherty, Franco-Watkins, and Thomas, 2008; Tversky and Kahneman, 1974), risk perception (Hertwig, Pachur, and Kurzenhäuser, 2005; Lichtenstein, Slovic, Fischhoff, Layman, and Combs, 1978), as well as factual judgment (Gigerenzer and Goldstein, 1996; Hertwig, Herzog, Schooler, and Reimer, 2008). In all these domains, the frequency of occurrence of a word in everyday language is positively correlated with the tendency to select that word as a response, and to evaluate the object denoted by that word as being large, important, truthful, or desirable. Indeed, fluency is believed to be one of the

⁵ Frame-based theories of coordination are sometimes divided into two broad classes, namely, team reasoning and level-*k* models. Theories of team reasoning assume that a group member follows the decision rule/frame that, if followed by other members, would be optimal for each of them (e.g., Crawford and Haller, 1990; Bacharach, 1993; Sugden, 1995; Casajus, 2000; Blume, 2000; Janssen, 2001; Alós-Ferrer and Kuzmics, 2013). By contrast, level-*k* theories assume a hierarchy of cognitive levels, whereby higher types best respond to lower-level players, anchoring their beliefs in the behavior of strategically naïve individuals (see Bacharach and Stahl, 2000, for a frame-based model of level-*k* reasoning). Relatedly, Charness and Sontuoso (2019) take a hybrid approach such that team reasoning is reduced to the case where one believes others play like oneself, given one’s partial awareness of frames.

mechanisms through which the *availability heuristic* operates (Tversky and Kahneman, 1973; Schwarz, Bless, Strack, Klumpp, Rittenauer-Schatka, and Simons, 1991).

The non-strategic literature above caused us to speculate that word frequency could be used as a proxy for prominence in strategic domains. For example, consider a pure coordination game where the strategy set is given by {paprika, curry, chili}.⁶ In this case, the strategy-label that is most frequently mentioned in everyday language may be most “fluent” (easy to process). Hence subjects could be drawn to that option, in the same way as they are drawn to fluent options in the non-strategic literature above. Yet – unlike the psychological literature – this paper aims to verify whether subjects *strategically* exploit word frequency. Here, one might argue that if there is common reason to believe that an option easily comes to mind to people with the same cultural background, then it may be optimal to select that option in pure coordination games.⁷ But what about games with unaligned incentives? And what about players with different cultural/linguistic backgrounds? To address these questions and test for the strategic use of word frequency, we propose the following three studies.

We designed Study 1 as a preliminary test for predicting behavior in 2-player (one-shot) pure coordination games with a finite set of strategies (played by culturally alike participants). A few points are worth noting. First, our strategy options present no *payoff asymmetries*; in fact, note that we only consider games without Pareto-rankable equilibria, which allows us to rule out a common driver of prominence as an explanation for our data patterns.⁸ Second, we designed each game by randomly drawing labels from lists of words in the same semantic domain (e.g., names of food ingredients): this means that our games have no obvious *frame asymmetries*, and hence we can rule out another common driver of prominence.⁹ Having found a strong correlation between word frequency (i.e., *n-gram* frequency computed in the general English

⁶ *Pure* coordination games are characterized by the following payoff structure: if players select the same strategy, they each receive an identical positive payoff (say, 1 currency unit); otherwise, they each receive nothing.

⁷ One mechanism supporting such a hypothesis (about the strategic use of word frequency) is that subjects may realize that their counterpart might be culturally alike and hence view the problem in a similar way. For evidence on “projection”, see Hedden and Zhang (2002), Sebanz, Knoblich, and Prinz (2003), and Rubinstein and Salant (2016).

⁸ For early evidence on coordination games with Pareto-rankable equilibria, see Cooper, DeJong, Forsythe, and Ross (1990); for the case of repeated games, see also Van Huyck, Battalio, and Beil (1990). For more recent experiments with Pareto-rankable equilibria, see Bardsley and Ule (2017) and Faillo, Smerilli, and Sugden (2017).

⁹ Recent experiments with *more or less* obvious frame asymmetries include, among others: Blume and Gneezy (2010); Bardsley, Mehta, Starmer, and Sugden (2010); Hargreaves Heap, Rojo Arjona, and Sugden (2014, 2017). Of particular interest is Hargreaves Heap et al.’s (2017) design, which elicits subjects’ beliefs about alternative heuristics that may drive behavior in coordination games. (Note that, unlike our studies, their design does not involve an *a-priori* measurable proxy for prominence, nor does it vary subjects’ incentives or cultural background.)

Google Books corpus) and behavior in coordination games, we designed two more studies to put to test our notion of prominence under different conditions.

Study 2 contrasts choice behavior in *(i)* a pure coordination game, against the behavior of participants in three alternative roles/conditions (with each condition featuring exactly the same list of labels). More specifically, we consider: *(ii)* the case in which a subject is prompted to *pick* an option, without any explicit objective; *(iii)* the case where a subject is prompted to avoid matching her counterpart's choice, under the assumption that her counterpart instead wants to match (i.e., the role of *Hider* in a Hide-and-Seek game); *(iv)* the case where a subject is prompted to match her counterpart's choice, under the assumption that her counterpart instead wants to avoid any such match (i.e., the role of *Seeker* in a Hide-and-Seek game). Since all our conditions involve the exact same options, note that if the effect of word frequency were merely due to an automatic (or naïve) response, then we should observe similar choice distributions across the four conditions. However, the data paint a different picture. Participants in problem *ii* (i.e., “Pickers”) were about as likely to select the most frequently-mentioned label as were participants in problem *i* (i.e., “Coordinators”). On the other hand, Hiders were less likely to select the most frequently-mentioned label than Seekers and, in turn, Seekers were less likely than Coordinators. As shall be discussed, this pattern indicates a boundedly rational, strategic use of word frequency that is consistent with a particular specification of level-*k* reasoning (Nagel, 1995; Stahl and Wilson, 1995; Costa-Gomes, Crawford, and Broseta, 2001).

We designed Study 3 to delve further into the strategic use of word frequency in pure coordination games. To that end, each of the games of Study 3 involves labels that we purposely selected so that the option with the highest word frequency differs between the American- and British-English vocabularies (as measured by the *n-gram* frequency in the American- and British-English Google Books corpora, respectively).¹⁰ We then varied the cultural/linguistic makeup of the subject pool by recruiting individuals residing in either the US or the UK; additionally, we manipulated our subjects' knowledge of the counterpart's country of residence. Consistent with our predictions, the data show that choice behavior differs between US and UK subjects and, in each case, it is positively related to the word frequency of the strategy-labels in the relevant vocabulary. Moreover, the data show that subjects are less likely to rely on word

¹⁰ For example, consider the game with strategy set {paprika, curry, chili}. There, “curry” has the highest *n-gram* frequency in *British English* and the lowest one in *American English*; conversely, “chili” has the highest *n-gram* frequency in *American English* and the lowest one in *British English*.

frequency as a means to guiding their behavior *if* they know that their counterpart resides in a different country. Put differently, if subjects are aware that their assigned partner is alike (in terms of cultural background), then they are more likely to select the label most frequently mentioned in their vocabulary.

We finally compared coordination rates that would be obtained if different subsamples were paired with each other, using Monte Carlo methods. In brief, successful coordination is more likely when subjects were *knowingly* paired with partners from their own country, as opposed to when they were knowingly or unknowingly paired with partners from a different country. Notably, subjects who were knowingly paired with partners from the same country exhibit expected coordination rates between 10 and 20 percentage points higher than chance.

To conclude, for the first time we propose and test an a-priori measurable proxy for prominence that explicitly rests on players' culture. Our experiments provide very robust evidence indicating that individuals play strategies fulfilling our notion of prominence in a (boundedly) rational manner. Remarkably, reliance on word frequency leads to higher rates of coordination than chance, and more so when individuals knowingly share a cultural background. The remainder of the article is organized in this manner: section II lays out the experimental procedures, sections III-V present our studies, and section VI concludes.

II - General Procedures

Our studies were conducted online between September and November 2016. A primary motivation for running online experiments is the ease with which the experimenter can control the cultural makeup of the subject pool. Another advantage is the ease with which the experimenter can vary subjects' knowledge of the fellow participants' cultural characteristics: this makes online experimentation optimal for testing culture-related hypotheses. (For a methodological discussion of extra-laboratory experiments, we refer the reader to Charness, Gneezy, and Kuhn, 2013.)

Participants were recruited through *Prolific Academic*, a crowdsourcing platform backed by Oxford University Innovation (<https://www.prolific.co/>). Participation in our study was limited to individuals with a Prolific Academic approval rate greater than 95%. At the beginning of each study, subjects were informed that they would be paired with a fellow participant at random, and that they would not know the identity of their counterpart or be able to

communicate with them. Participants' responses were incentivized, as specified in the following sections. No subject was allowed to participate in more than one study.

As a proxy for the prominence of strategy-labels, we used the *case-insensitive average yearly n-gram frequency* (henceforth simply NGRAM) of the corresponding words in the Google Books corpus, for books published after 2000. NGRAM values represent the fraction of times in which phrases (in our case, words) occur in the corpus of interest; for Studies 1, 2 we used the general English corpus, whereas for Study 3 we used the American-English and British-English corpora. The values were obtained through the Google n-gram tool in August 2016, shortly before running the studies (<https://books.google.com/ngrams/info>). Note that n-gram frequencies are a reliable, standard metric for word popularity in corpus linguistics (Michel et al., 2011; Garg, Schiebinger, Jurafsky, and Zou, 2018; Hills, Proto, Sgroi, and Seresinhe, 2019).

III - Study 1

Demographics. The subject pool for Study 1 consisted of 91 US resident individuals. The average participant was 33 years old, and 57% of the subjects were male. Participants took less than 10 minutes to review the instructions and complete all the tasks; they received a 0.5 GBP participation fee (in addition to the payoffs earned in each game), which is on par with typical wages on Prolific Academic or other Internet marketplaces such as MTurk.

Design. We designed Study 1 as a preliminary test for predicting behavior in 2-player pure coordination games. This study involves a series of (one-shot) games, with each game featuring a 3-element strategy set, such that: each member of a pair receives GBP 0.10 if both players choose the same option; each member of a pair receives nothing otherwise. Figure 1 represents the game structure in bimatrix form (there, for expositional purposes the set of strategies is denoted by {X, Y, Z}; note that subjects were not provided any such figure). Subjects played 10 instances of the game, with each instance differing from the others only in the names of the three options. Each subject was assigned the same partner for all the (10) games. *No feedback* was provided between games.

We ran two versions of the study: Version A's options consist of names of countries, whereas Version B's options consist of names of food ingredients; both versions are shown in Table 1 below. The reason we designed two versions is to verify that the (presumed) prominence of frequently-mentioned labels does not depend on the characteristics of a specific

collection. Note that the option sets for Version A were obtained by drawing member states of the United Nations at random. The option sets for Version B were obtained by drawing words at random from a list of ingredients scraped off the Epicurious cooking website (<https://www.epicurious.com/>).¹¹

		Player 2		
		X	Y	Z
Player 1	X	0.10	0	0
	Y	0.10	0	0
	Z	0	0.10	0

Figure 1 - The coordination game. The bottom-left and top-right numbers in each cell represent the monetary payoffs to Player 1 and Player 2, respectively. (For expositional purposes, the set of strategies is denoted by {X, Y, Z}; subjects were not provided any such figure.)

Participants were randomly assigned to the two versions: 48 subjects were assigned to Version A, and 43 subjects to Version B.¹² Average earnings were GBP 0.63 and 0.52 for Version A and B, respectively (not including subjects’ participation fees). Before presenting the experimental results, we shall note that the order of the games was randomized across subjects. By contrast, the order of the three options – in a given game – was determined prior to the experiment at random, and was *identical* across subjects. More precisely, in each game the strategy-labels were arranged in a column, with options X, Y, and Z of Table 1 below being respectively displayed at the top, center, and bottom of the list. For the experimental instructions and screenshots, see *Appendix B*.

¹¹ It is worth clarifying a general point in relation to our option sets. It is possible that – around the time our study was conducted – some options in our lists were being mentioned in the news or social media more often than usual, a fact that would not be immediately accounted for by the n-gram frequency at the time. Nevertheless, we believe that our *randomized* selection of labels overall controls for any such “random shocks”.

¹² Given the odd number of subjects in Version B, one participant was assigned two partners (but received compensation for playing with either one, at random); the two partners were treated like any other participant.

Version A			Version B				
	[option X]	[option Y]	[option Z]		[option X]	[option Y]	[option Z]
1	Kyrgyzstan .4435	Tuvalu .1139	Morocco 2.4424 (h)	1	mandarin .7727	soybean .8652	cinnamon 1.3619 (h)
2	Turkey 2.1062 (h)	Jamaica .6537	Senegal .2399	2	raspberry 2.0997 (h)	sauerkraut .6453	scallion .2549
3	Yemen .6588	Benin .4876	Jamaica 1.8535 (h)	3	cashew .2170	vanilla 1.5978 (h)	yogurt 1.1850
4	Spain 2.2128 (h)	Angola .1943	Norway .5927	4	meat 2.5905 (h)	cocoa .3052	raspberry .1041
5	Bolivia 2.1420 (h)	Kyrgyzstan .5533	Palau .3046	5	horseradish 1.0992	rhubarb 1.1067 (h)	tarragon .7940
6	Afghanistan .3663	Germany 2.4533 (h)	Ghana .1802	6	butter 2.3233 (h)	peppermint .1339	cocoa .5426
7	Nepal .9101	Uzbekistan .3106	Thailand 1.7791 (h)	7	pineapple 1.6064 (h)	nutmeg .9394	buckwheat .4540
8	Bahamas 1.1349	Botswana 1.7162 (h)	Tuvalu .1487	8	apple .9075	milk 1.9335 (h)	tuna .1588
9	Bahamas .6005	Eritrea .4705	Ghana 1.9289 (h)	9	buckwheat .1740	tomato 1.9814 (h)	citrus .8445
10	Ukraine .7379	Jordan 1.9148 (h)	Zambia .3472	10	brandy .9506	pumpkin .5337	tomato 1.5156 (h)

Table 1 - The option sets for Study 1. The left and right panels refer to Version A and B, respectively. Below each strategy-label is the *relative n-gram frequency* of that label, computed from the general English corpus including books published after the year 2000. Note: for each option, the reported number is obtained by dividing the NGRAM value of its label by the mean of the values of the three labels in the game (to normalize the data, simply divide each value by 3). For visual clarity, the option with the relatively highest NGRAM value is marked with an “h”.

Results. The average participant in Study 1 chose the strategy associated with the highest, middling and lowest NGRAM value respectively 45.27%, 39.34% and 15.39% of the time (specifically, subjects selected the option with the highest NGRAM 45.83% of the time in Version A, and 44.65% of the time in Version B on average; for a bar graph of the distributions

of individual-level choices in each of the games, see *Appendix A*). These data patterns are clearly suggestive of a positive impact of the labels' n-gram frequency on strategic play.¹³ The following tests provide extensive evidence in support of the strategic use of word frequency, while addressing a potential confound.

One may argue that labels that are displayed in a particular position might be perceived as more salient by some subjects. Thus, we now verify that the impact of word frequency is not confounded by the position of a label on the screen (i.e., top, center, or bottom of the list). To do so, we shall compute the frequency with which a subject chooses the n -th position, in games where the n -th label *does* and *does not* have the highest NGRAM value.

We start by considering the label displayed first (i.e., at the top of the list). This exercise reveals that when the top label has the highest NGRAM value, it is chosen 55.03% of the time on average; by contrast, the top label is chosen 45.36% of the time whenever it does not have the highest NGRAM value. A two-tailed Wilcoxon sign-ranked test shows that the difference is strongly significant ($N = 91$ obs., $z = 3.042$, $p = 0.002$).¹⁴ Performing the same analysis with respect to the other positions corroborates the trend. In fact, when the center label has the highest NGRAM value it is chosen 36.44% of the time, compared with 21.08% whenever it does not have the highest value ($N = 91$ obs., $z = 5.305$, $p = 0.000$). Furthermore, when the bottom label has the highest NGRAM value it is chosen 43.40% of the time, compared with 15.75% whenever it does not have the highest value ($N = 91$ obs., $z = 6.421$, $p = 0.000$). In summary, *the n -th strategy option is chosen significantly more often when its label has the highest NGRAM value, compared with when it does not have the highest value.*

Moving on, we note that since the three options (in any of our games) constitute symmetric strategies, here Harsanyi and Selten (1988) would argue that the rational solution is to play a “symmetry-invariant” equilibrium, assigning each option the same probability. In their

¹³ Interestingly, these summary data are comparable to the behavioral patterns observed in games where one of the strategies is devised by the experimenter as the *obvious* “*odd-one-out*”. For example, Mehta, Starmer, and Sugden (1994) designed four coordination games where the strategy options are represented by visual objects (i.e., questions no. 17-20, p. 669). In two of those games the odd-one-out was chosen by 44.4% of participants, whereas in the other two games it was chosen by about 64% of participants; in this regard, note that the Mehta et al. (1994) games feature a 15-element strategy set, which renders the odd-one-out more prominent (recognizably different) than it would have been in a 3-element strategy set such as ours.

¹⁴ The test uses one observation per subject, consisting of the difference between the two above-described rates (i.e., the frequency with which a subject plays top, in games where the top label has the highest NGRAM value *and* in games where the top label does not have the highest value). Note that the Wilcoxon sign-ranked test is the non-parametric analog to the paired samples t-test.

view, such a solution has the benefit of being unique and it ensures that a renaming of the strategies *cannot* ever affect game play (Harsanyi and Selten, 1988, pp. 70-74.) In this case, their proposed solution is not supported by the data, as our choice distribution significantly differs from the fully mixed equilibrium strategy profile that assigns equal probability to each strategy ($N = 91$ obs., $T^2 = 127.67$, $p = 0.000$ under a Hotelling's T-squared generalized means test, conducted on the *sample of per-subject mean choices*; note that the Hotelling's test is simply a multivariate generalization of the t-test).¹⁵ Relatedly, we stress that the fully mixed equilibrium implies a coordination rate of 0.33, whereas the expected coordination rate resulting from our sample is roughly 0.50 on average (specifically, 0.57 for Version A and 0.49 for Version B; this means that the payoffs earned by participants in Version A and B are respectively 72% and 48% higher than the payoffs subjects would earn by coordinating on the fully mixed equilibrium).^{16,17}

In concluding, we note that we designed Study 1 as a preliminary test for predicting behavior in coordination games. The method of analysis employed so far has involved mean observations, thereby discarding a fair amount of information. To shed light on the strategic use of word frequency, below we consider some between-subjects designs; we then perform a more sophisticated analysis in such a way to account for the characteristics of each triplet of labels.

IV - Study 2

Demographics. The subject pool for Study 2 consisted of 160 US resident individuals. The average participant was 30 years old, and 58% of the subjects were male. As with Study 1, participants took less than 10 minutes to review the instructions and complete all the tasks; they received a 0.5 GBP participation fee, in addition to the payoffs earned in each game (if any), as specified below.

¹⁵ The sample of (per-subject) mean observations is obtained as follows. First, for each choice of subject i , assign a value of 1 or 0 to indicate if the option with the *highest* NGRAM was chosen or not; then, take the average across all the games played by subject i . Similarly, assign a value of 1 or 0 to indicate if the option with the *middling* NGRAM was chosen or not, and take the average across games. The same applies to the option with the *lowest* NGRAM.

¹⁶ In keeping with previous studies, we report expected coordination rates (as opposed to actual frequencies of coordination; see for example Mehta et al., 1994, and Crawford et al., 2008). In fact, actual frequencies of coordination are affected by the eventual pairing of partners, thereby resulting in a biased metric in smaller samples.

¹⁷ The expected coordination rate gives the probability that two randomly drawn subjects choose the same strategy in a randomly selected game. For each version of Study 1, we calculate this rate using Monte Carlo methods. The pseudo code is as follows: (A) pick two participants at random; (B) pick one of the 10 games at random; (C) if both participants chose the same strategy consider it an instance of successful coordination, otherwise consider it unsuccessful; (D) repeat steps A-C 100,000 times, then calculate the relative frequency of successful coordination.

Design. The objective of the study is to test for the strategic use of word frequency by systematically varying players' incentives. To that end, we designed a few variants of the simple coordination game of Study 1, in such a way to incentivize or disincentivize coordination for either player. In order to check for replicability, Study 2 features the exact same option triplets as in Version A of Study 1 (see the left panel of Table 1 above).¹⁸ That being said, Study 2 involves the following four roles/conditions.

- a. **Coordinate:** This is a replication of Version A of Study 1 that, among other purposes, serves to verify the robustness of our earlier results. Participants in this condition (“Coordinators”) were paired with other participants in the same condition, and were so informed.
- b. **Pick:** Participants in this condition (“Pickers”) were presented with the same labels as in the Coordinate condition, except that they were asked to merely pick one of the three given options. That is, participants were not assigned a partner, nor did they receive any additional payoffs on the basis of their choices; hence, they had no strategic incentive to select one label over another.
- c. **Seek:** This condition features the same strategy-labels as in the Coordinate condition, except that the incentive structure reflects the role of “Seeker” in the Hide-and-Seek game—Figure 2 below represents the payoff structure of this game. As can be seen there, a Seeker receives GBP 0.10 if both members of a pair choose the same strategy, and nothing otherwise. Participants in this condition were paired with participants in the Hide condition below, and were so informed.
- d. **Hide:** Again, this condition features the same strategy-labels as in the Coordinate condition, but the incentive structure reflects the role of “Hider” in the Hide-and-Seek game. As can be seen in Figure 2, a Hider receives GBP 0.10 if members of a pair choose different strategies, and nothing otherwise. Participants in this condition were paired with participants in the Seek condition above, and were so informed.

Each subject completed 10 tasks in the same role/condition, with each task differing from the others only in the names of the three options (see the left panel of Table 1 above). In the Coordinate, Seek, and Hide conditions each subject was assigned the same partner for all the (10) games, and was so informed.¹⁹ *No feedback* was provided between games.

¹⁸ Note that here we focused on Version A – i.e., names of countries – simply to be economical; in Study 3 below we will resume investigating the impact of word frequency in the context of food ingredients.

¹⁹ Given that in the *Hide* condition there were less participants than in the *Seek* condition, (for the mere purpose of calculating the payoffs of the extra Seekers) nine Hiders were matched with two Seekers, but received compensation for playing with either one at random.

Before discussing our predictions, we note that (as with Study 1) the order of the tasks/games was randomized across subjects. On the other hand, the order of the three options – in a given task – was determined prior to the experiment at random, and was identical across subjects. Specifically, the order of the three options was the same as in Study 1; i.e., the labels were arranged in a column, with options X, Y, and Z of Table 1 (left panel) being respectively displayed at the top, center, and bottom of the list. For the experimental instructions and screenshots, see *Appendix B*.

		Seeker		
		X	Y	Z
Hider	X	0.10 0	0 0.10	0 0.10
	Y	0 0.10	0.10 0	0 0.10
	Z	0 0.10	0 0.10	0.10 0

Figure 2 - The Hide-and-Seek game structure in bimatrix form. The bottom-left and top-right numbers in each cell represent the monetary payoffs to the *Hider* and *Seeker*, respectively.

Since all our conditions involve the exact same labels, if the effect of word frequency were merely due to an automatic (or naïve) response, then we should observe similar choice distributions across conditions. If instead participants used word frequency in a strategic manner, then frequently-mentioned labels should be more (less) likely selected when there are incentives to match (mismatch) others, with the magnitude of the changes varying with subjects' strategic sophistication. Below we model our intuition with a particular specification of level- k reasoning (Nagel, 1995; Stahl and Wilson, 1995; Costa-Gomes, Crawford, and Broseta, 2001; Camerer, Ho, and Chong, 2004; Costa-Gomes and Crawford, 2006).

Level- k theories posit a hierarchy of player types defined by the level of sophistication with which each player reasons. Specifically, level- k types anchor their beliefs in a non-strategic $L0$ type and adjust them via iterated best responses, so that $L1$ players best respond to $L0$ players,

$L2$ best respond to $L1$, and so on.²⁰ In what follows, we implement this approach by formulating a set of assumptions that are relevant for our setting; this exercise will formally generate our predictions. In short, we shall assume that:

- (i) A non-strategic $L0$ type in the *Coordinate*, *Seek*, or *Hide* conditions behaves like a participant in the *Pick* condition;
- (ii) Players at levels *above* $L0$ believe that the distribution q of $L0$ choices has a peak at the option with the highest NGRAM value;
- (iii) There are no players at $L3$ or higher.

A few comments are in order. We note that level- k theories commonly assume that $L0$ types do not engage in strategic reasoning, but simply randomize between options according to some probability distribution q . In particular, Crawford and Iriberri (2007) assume that such a distribution is non-uniform – positing that $L0$ types are relatively more likely to select salient labels, compared with other labels – without however defining “salience” in general terms.

For the purposes of generating hypotheses, we partly sidestep the issue of specifying the probability distribution q by defining it empirically (on the basis of Pickers’ behavior), as per assumption (i). Then, since our subjects obviously do not observe that distribution, we posit that players at levels *above* $L0$ believe that “ $L0$ types are more likely to select the option with the highest NGRAM than to select any other option”, as per assumption (ii). Lastly, we note that (iii) is a simplifying assumption, based on previous empirical evidence about subjects’ strategic sophistication. For example, Arad and Rubinstein (2012) observe that level- k experiments have shown that «the most frequent types are usually $L1$ and $L2$, whereas higher-level types are rare» (p. 3561). For a similar point, see also Penczynski’s (2016) analysis of Hide-and-Seek games.

What does the above entail in terms of behavioral predictions? As usual, $L1$ types will best respond to (their beliefs about) $L0$ behavior; $L2$ types will accordingly adjust their beliefs, and best respond to $L1$ types. Specifically, $L1$ Coordinators will best respond to their beliefs about $L0$ behavior (which in our case are defined by assumption ii), and therefore will select the option with the highest NGRAM with probability one. Then, $L2$ Coordinators will best respond

²⁰ Note that some applications of level- k reasoning differ in their assumptions as to whether there are actually any players at $L0$ (as will be clear, such assumptions do not qualitatively affect our predictions). Applications further differ in their assumptions as to the players’ randomizing behavior at $L0$. Another element in regard to which models differ is the players’ depth of reasoning about other types; in particular, some models assume that players at each level *above* $L0$ best respond to a probability mix of the decisions of all levels below their own, as opposed to best responding to the one level immediately below. For discussion, see Crawford, Costa-Gomes, and Iriberri (2013) and Mauersberger and Nagel (2018).

to *L1* Coordinators, thereby selecting the option with the highest NGRAM (with probability one). Moving on, we note that Seekers wish to select the label believed to be the modal choice of Hiders at the level below, whereas Hiders wish to avoid the label believed to be the modal choice of Seekers at the level below: here, this implies that the option with the highest NGRAM will be respectively selected by *L1* Hiders and *L1* Seekers with probability zero and one; hence, the option with the highest NGRAM will be selected by both *L2* Hiders and *L2* Seekers with probability zero.

In summary, the experiment aims to verify if subjects' behavior is compatible with a strategic use of word frequency. If it were not, we should observe the same choice distributions across roles/conditions. If instead subjects used word frequency in a (boundedly) rational manner, then – based on the assumptions above – behavior should vary across roles as follows.

H1: *Coordinators select the most frequently-mentioned label as often as (or more often than) Pickers.*

H2: *Coordinators select the most frequently-mentioned label more often than Seekers.*

H3: *Seekers (and Pickers) select the most frequently-mentioned label more often than Hiders.*

Formally, the assumptions above imply that the “Coordinators’ probability of selecting the most frequently-mentioned label” (i.e., the option with the highest NGRAM) is defined by $p_{COORD} = (l_0 \cdot q_H) + (l_1 \cdot 1) + (l_2 \cdot 1)$, where l_k denotes the share of L_k players in our subject pool while q_H denotes the frequency with which L_0 types select the option with the highest NGRAM (as defined by assumption i). Thus, a term in the expression above represents the probability that the relevant share of L_k Coordinators in our subject pool select the option with the highest NGRAM. (Incidentally, we stress that L_0 behavior by definition is the same across roles.) Next, the “Seekers’ probability of selecting the option with the highest NGRAM” is given by $p_{SEEK} = (l_0 \cdot q_H) + (l_1 \cdot 1) + (l_2 \cdot 0)$, where the second and third terms refer to *L1* and *L2* Seekers, respectively. Further, the “Hiders’ probability of selecting the option with the highest NGRAM” is given by $p_{HIDE} = (l_0 \cdot q_H) + (l_1 \cdot 0) + (l_2 \cdot 0)$, where the second and third terms respectively refer to *L1* and *L2* Hiders. Now, assuming that there are strategic players in our subject pool (i.e., $l_1, l_2 > 0$, with $l_0, l_1, l_2 \in (0,1)$ and $l_0 + l_1 + l_2 = 1$), then – under the standard assumption that the distribution of types is the same across roles/conditions – the above entails that $p_{COORD} > p_{SEEK} > p_{HIDE}$ and $p_{COORD} \geq p_{PICK} \geq p_{HIDE}$, where p_{PICK} denotes the “Pickers’

probability of selecting the option with the highest NGRAM”, with $p_{PICK} \equiv q_H$ as per assumption (i).²¹ Conversely, the above implies that the likelihood of choosing the option with the lowest NGRAM will possibly rise when moving from Coordinate (or Pick) to Seek to Hide.

Results. Table 2 presents mean choices in each of the four roles/conditions, given a classification of the strategy-labels based on their relative n-gram frequency, as per Table 1 above. (For a bar graph of the distributions of individual-level choices in each of the games and conditions, see Appendix A.) By taking a glance at Table 2, the reader will notice that the mean distribution of choices varies with each condition. In fact, the *most* frequently-mentioned label (i.e., the option with the highest NGRAM) was chosen less and less often when moving from Coordinate (or Pick) to Seek to Hide. Consequently, the *least* frequently-mentioned label (i.e., the option with the lowest NGRAM) was chosen more and more often when moving from Coordinate (5.71%) to Pick (13.33%) to Seek (22.27%) to Hide (30.29%). Taken together, these patterns seem to confirm that subjects used word frequency in a (boundedly) rational manner, whereby the higher the word frequency of a label, the lower the likelihood of choosing that option when moving from Coordinate to Hide. Put differently, the lower the word frequency of a label, the higher the likelihood of choosing that option when moving from Coordinate to Hide. In what follows we further examine these trends.

We begin by reporting a Kruskal-Wallis test, which confirms significant differences in the choice of the option with the *highest* NGRAM across conditions ($N = 160$ obs., $\chi^2_3 = 10.477$, $p = 0.014$, two-tailed; note that in order to satisfy the assumption of independence of observations, all our non-parametric tests are conducted on the sample of per-subject mean choices, as described in footnote 15, p. 11). Similarly, the same test confirms significant differences in the choice of the option with the *lowest* NGRAM across conditions ($N = 160$ obs., $\chi^2_3 = 34.353$, $p = 0.000$, two-tailed).^{22,23}

²¹ The weak inequalities are due to the fact that the modeler has no a-priori knowledge of q_H , with $q_H \in [0,1]$. On a different note, we stress that since our focus is on behavioral comparisons across roles (as opposed to identifying the empirical distribution of levels), we need not make any further assumptions in order to generate our hypotheses.

²² Also, Hotelling’s T-squared generalized means tests (conducted on the samples of per-subject mean choices) reveal that the distribution of choices in each of the Coordinate, Pick, and Seek conditions differs from the fully mixed equilibrium assigning equal probability to all strategies (for *Coordinate*: $N = 42$ obs., $T^2 = 389.71$, $p = 0.000$; for *Pick*: $N = 39$ obs., $T^2 = 55.72$, $p = 0.000$; for *Seek*: $N = 44$ obs., $T^2 = 20.66$, $p = 0.000$). However, the same test shows that the distribution of choices in the *Hide* condition does not differ from the fully mixed equilibrium, a fact that might be interpreted as indirect evidence in support of H3 (we shall test that hypothesis below).

	<i>Coordinate</i>	<i>Pick</i>	<i>Seek</i>	<i>Hide</i>
Choice by word frequency				
Strategy-label with <i>highest</i> NGRAM metric is chosen [f_H], %	46.19 (.1464)	48.46 (.1646)	43.64 (.1556)	36.86 (.1548)
Strategy-label with <i>middling</i> NGRAM metric is chosen [f_M], %	48.10 (.1596)	38.21 (.1211)	34.09 (.1661)	32.85 (.1202)
Strategy-label with <i>lowest</i> NGRAM metric is chosen [f_L], %	5.71 (.0914)	13.33 (.1675)	22.27 (.2044)	30.29 (.2121)
Total, %	100	100	100	100
Total # triplets (1,600)	420	390	440	350
Subjects (160)	42	39	44	35

Table 2 - (Per-subject) mean choices, given a classification of the labels based on their relative n-gram frequency; in parentheses is the standard deviation. Note: the number of triplets is obtained by multiplying the number of tasks/games (i.e., 10) by the number of participants in each role/condition.

In order to address our hypotheses, later on we report a formal econometric analysis that accounts for the characteristics of each triplet of labels, such as the position of the labels on the screen (i.e., top, center, or bottom of the list) and their numerical NGRAM values. Before doing so – to sketch a rough outline of the main patterns emerging from our dataset – we shall present some non-parametric tests conducted on the sample of per-subject mean observations (i.e., the tests use one observation per participant). We start by comparing behavior in the *Coordinate* and *Pick* conditions: a one-tailed test allows us to check the “alternative hypothesis” that the most frequently-mentioned label (i.e., the option with the highest NGRAM) is selected strictly less often in *Coordinate* than in *Pick*.²⁴ In short, a one-tailed Wilcoxon-Mann-Whitney test shows no evidence of a significant decrease in the choice of the *most* frequently-mentioned label when

²³ Incidentally, we note that the expected coordination rate resulting from our sample is 0.44 (for subjects in a *Coordinator* role). This implies that, on average, the payoff to our Coordinators is 33% higher than the payoff subjects would obtain by playing the fully mixed equilibrium. In the case of the Hide-and-Seek game, the average payoff to subjects in a *Seeker* role is about 2% higher than the payoff that would be obtained by a hypothetical Seeker who randomizes uniformly over labels; then, the average payoff to subjects in a *Hider* role is about 1% lower than the payoff that would be obtained by a hypothetical Hider who randomizes uniformly over labels.

²⁴ Note that H1 says that Coordinators select the most frequently-mentioned label as often as (or more often than) Pickers. Because of the weak inequality, here we shall test against the alternative hypothesis that the most frequently-mentioned label is selected strictly less often in *Coordinate* than in *Pick*.

moving from Pick to Coordinate ($N = 81$ obs., $Z = -1.008$, $p = 0.1567$). Conversely, a one-tailed Wilcoxon-Mann-Whitney test provides evidence of a significant decrease in the choice of the *least* frequently-mentioned label when moving from Pick to Coordinate ($N = 81$ obs., $Z = -2.726$, $p = 0.003$). These tests provide some very preliminary evidence in support of H1.

A similar analysis shows no significant difference between the *Coordinate* and *Seek* conditions, with respect to the choice of the *most* frequently-mentioned label. Yet, a one-tailed Wilcoxon-Mann-Whitney test provides evidence of a significant increase in the choice of the *least* frequently-mentioned label when moving from Coordinate to Seek ($N = 86$ obs., $Z = -4.235$, $p = 0.000$). The latter result might be viewed as indirect evidence in support of H2, which warrants further testing: later on, the econometric analysis will shed light on these data patterns.

Non-parametric tests additionally show that the *most* frequently-mentioned label was selected more often in *Pick* than in *Hide*, providing some preliminary support for H3 ($N = 74$ obs., $Z = 3.101$, $p = 0.000$, one-tailed Wilcoxon-Mann-Whitney test). Similarly, the *most* frequently-mentioned label was selected significantly more often in *Seek* than in *Hide* ($N = 79$ obs., $Z = 1.978$, $p = 0.023$, one-tailed Wilcoxon-Mann-Whitney test), which again supports H3. (Unsurprisingly, the same test shows that the *least* frequently-mentioned label was selected significantly more often in *Hide* than in *Pick*, and significantly more often in *Hide* than in *Seek*.)

In summary, our non-parametric tests confirm a trend where the higher the word frequency of a label, the lower the likelihood of choosing that option when moving from Coordinate (or Pick) to Hide. Put differently, there is a trend where the lower the word frequency of a label, the higher the likelihood of choosing that option when moving from Coordinate to Pick to Seek to Hide. That being said, we note that since the above tests are conducted on the sample of (per-subject) mean choices, they do not account for differences in individual responses across tasks/games, or for differences in the numerical NGRAM values; also, those tests do not control for the labels' position on the screen.

For the reasons above, we shall now corroborate our analysis by reporting the results of alternative-specific conditional logit models (“asclogit”; i.e., McFadden’s choice model, 1973). These models will provide the ultimate test of our hypotheses, controlling for the characteristics of each game (i.e., triplet of labels) to which each subject was exposed. Below we report the main findings while we refer the reader to Appendix A for the full econometric tables.

Overall, the models confirm a significant positive effect of word frequency on choice behavior, across treatments. In particular, the higher the n-gram frequency of a label, the more

likely it is for that strategy option to be selected (regardless of the label's position on the screen). Moreover, when comparing behavior in *Coordinate* and *Pick*, the model indicates no significant difference in the relative impact of word frequency between these conditions (see variable *WFC* in model [1] of Table 1A, in Appendix A). This implies that there is no difference in the probability of choosing the most frequently-mentioned label between *Coordinate* and *Pick*, which confirms the previous evidence in regards to H1.

Next, when contrasting behavior in *Coordinate* against *Seek*, we find a significant difference in the relative impact of word frequency on choice, controlling for the labels' position on the screen. That is, a label with a higher NGRAM value is more likely to drive the choices of Coordinators than Seekers (coef. = $-.119$, $z = -2.44$, $p = 0.015$, two-tailed asclogit conducted on the sample of individual observations, with standard errors adjusted for clustering on 86 subjects; see variable *WFC* in model [2] of Table 1A). The result evidently supports H2.

Also, contrasting behavior in *Pick* against *Hide*, we find a significant difference in the relative impact of word frequency on choice (controlling for the labels' position on the screen); that is, a label with a higher NGRAM value is more likely to drive the choices of Pickers than Hiders (coef. = $-.201$, $z = -3.77$, $p = 0.000$, two-tailed asclogit conducted on the sample of individual observations, with standard errors clustered on 74 subjects; see variable *WFC* in model [3] of Table 2A). Finally, when contrasting behavior in *Seek* against *Hide*, we find again a significant difference in the relative impact of word frequency on choice (controlling for the labels' position on the screen): a label with a higher NGRAM value is more likely to drive the choices of Seekers than Hiders (coef. = $-.2145$, $z = -2.14$, $p = 0.033$, two-tailed asclogit conducted on the sample of individual observations, with standard errors clustered on 79 subjects; see variable *WFC* in model [4] of Table 2A). The results support H3.

To conclude, the data provide strong support for our hypotheses (please refer to Appendix A for an extended commentary on the econometric analysis). Despite the fact that our conditions involve the same option triplets, more frequently-mentioned labels were selected less often when moving from *Coordinate* to *Seek*, from *Pick* to *Hide*, and from *Seek* to *Hide*. These data patterns confirm that individuals select strategies that fulfill our characterization of prominence, and they do so in a (boundedly) rational manner consistent with our level- k specification. In the remainder of the article we shall delve into the strategic use of labels in pure coordination games.

V - Study 3

Demographics. The subject pool for Study 3 consisted of 80 individuals, of which half were US residents and half were UK residents. Participants were recruited through the Prolific Academic platform, at the same time. In the US-residents sample, the average participant was 29 years old, and 75% of the subjects were male. In the UK-residents sample, the average participant was 31 years old, and 63% of the subjects were male. As with our previous studies, participants took less than 10 minutes to review the instructions and complete all the tasks; they received a 0.5 GBP participation fee, in addition to the payoffs earned in each game, as specified below.

Design. In what follows we describe a set of treatments that are intended to further illuminate the strategic use of labels in coordination games. The present study involves exactly the same pure coordination structure as in Study 1 (see Figure 1 above), except that this time our games feature different labels than before. Unlike our previous studies – which featured randomly generated labels – in this case we purposely selected triplets of labels so that the option with the highest word frequency differs between the American- and British-English vocabularies (as measured by the relevant NGRAM in the American- and British-English Google Books corpora).

Specifically, subjects played 10 instances of the coordination game in Figure 1, with each instance differing from the others only in the names of the three options (displayed below). Each subject was assigned the same partner for all the (10) games, and was so informed. *No feedback* was provided between games. The order of the games was randomized across subjects. Instead, the order of the three options – in a given game – was determined prior to the experiment at random, and was identical across subjects. That is, in each game the strategy-labels were arranged in a column, with options X, Y, and Z of Table 3 being respectively displayed at the top, center, and bottom of the list. For the experimental instructions and screenshots, see *Appendix B*.

As we previously noted, a key feature of this experiment is that we varied the cultural makeup of the subject pool (by recruiting samples of US residents and of UK residents). Additionally, we manipulated the subjects' perception of their counterpart's cultural/linguistic background: we did so by providing subjects with different information as to their partner's country of residence. More precisely, each participant – whether in the US or in the UK sample – was assigned to one of the following “information conditions”.

- a. **NO-Info**: Participants in this condition received no information about their partner’s country of residence (hence, this condition is identical to Study 1, except that we used different strategy-labels). The instructions in this condition stated: *“Please choose one option. Each of you and your partner receive £0.10 if you both choose the same option, £0 otherwise”*.
- b. **Know-UK (“k-UK”)**: Participants in this condition were told that their partner resided in the UK. Specifically, subjects were shown the following message: *“Please choose one option. Each of you and your partner receive £0.10 if you both choose the same option, £0 otherwise. Your partner is a Prolific worker who resides in the UK. Your partner may or may not know where you reside”*.
- c. **Know-US (“k-US”)**: Participants in this condition were told that their partner resided in the US. Specifically, subjects were shown the following message: *“Please choose one option. Each of you and your partner receive £0.10 if you both choose the same option, £0 otherwise. Your partner is a Prolific worker who resides in the US. Your partner may or may not know where you reside”*.²⁵

In summary, we recruited samples of US and UK residents; each subject was then randomly assigned to one of the three information conditions, irrespective of the subject’s own residence.

We stress that if the effect of word frequency in coordination games were due to an automatic (or naïve) response, then we should observe similar choice distributions across conditions and countries. If however subjects used word frequency in a strategic manner, then we should find that options with higher NGRAM values (in the subjects’ respective vocabularies) are more often selected when subjects think that their partner is culturally alike. In that case, subjects would realize that culturally alike people may view the problem similarly. That is, if a label comes to mind easily to a subject, then she will realize that it may also come to mind easily to others with the same vocabulary. This leads to the following hypotheses.

H4: *Choice behavior differs between US and UK residents, and it is positively related to the labels’ n-gram frequency in the vocabulary of the respective countries.*

H5: *Subjects who are informed that their partner resides in the same country select the most frequently-mentioned label (in their vocabulary) more often than subjects who are unaware of their partner’s residence. In turn, the latter select the most frequently-mentioned label more often than subjects who are informed that their partner resides in a different country.*

²⁵ The reason we tell subjects that their partner may or may not know where they reside is the following. For the purpose of calculating coordination rates via Monte Carlo methods, such a wording permits us to virtually match each subject with any participants in the relevant country (regardless of those participants’ information condition).

	[option X]	[option Y]	[option Z]		[option X]	[option Y]	[option Z]
1	paprika	curry	chili	6	sardines	tuna	cod
US	.5718	.3859	.6963 (h)	US	.4648	.6225 (h)	.4439
UK	.4281	.6140 (h)	.3036	UK	.5351	.3774	.5560 (h)
2	Bordeaux	Chianti	Syrah	7	cheesecake	scones	tiramisu
US	.4784	.6574 (h)	.5222	US	.6852 (h)	.4438	.5849
UK	.5215 (h)	.3425	.4777	UK	.3147	.5561 (h)	.4150
3	venison	lamb	pork	8	burrito	panini	kebab
US	.5240	.4543	.5816 (h)	US	.7599 (h)	.3145	.3579
UK	.4759	.5456 (h)	.4183	UK	.2400	.6854 (h)	.6420
4	peach	pineapple	pear	9	parsnip	beetroot	shallot
US	.5110	.5731 (h)	.4359	US	.5492	.2788	.7001 (h)
UK	.4889	.4268	.5640 (h)	UK	.4507	.7211 (h)	.2998
5	blueberry	blackberry	gooseberry	10	oatmeal	porridge	granola
US	.6741 (h)	.4838	.2700	US	.5838	.3529	.8263 (h)
UK	.3258	.5161	.7299 (h)	UK	.4161	.6470 (h)	.1736

Table 3 - The option sets for Study 3. The first and second row below the strategy-labels respectively report the *American-English* and *British-English* relative n-gram frequency of the corresponding word, for books published after the year 2000. For visual clarity, the option with the relatively highest NGRAM value is marked with an “*h*”.²⁶

We note that the hypotheses above are consistent with a level-*k* specification allowing for multiple (alternative) *L0* types, whereby the modal choice of each *L0* type corresponds to the option with the highest NGRAM value in that type’s vocabulary. Given this, players at *L1* best respond to a convex combination of their beliefs about each *L0* type; as usual, players at *L2* then best respond to (their beliefs about) *L1* behavior, and so on.

Results. Table 4 presents mean choices, given a classification of the strategy-labels based on their n-gram frequency as follows. The left panel provides summary data by pooling all the

²⁶ For any given label, the number reported in the first row (“US”) below the label is obtained by performing the following operations: (i) divide the American-English NGRAM value of the label by the sum of the values of the three labels in the game; (ii) perform the same operation as before, except this time use British-English NGRAM values; (iii) divide the outcome of (i) by the sum of the outcomes of (i) and (ii). The rationale behind step (iii) is to ensure that our measure of word frequency is a function of both vocabularies and hence applies to all the subjects, regardless of their information condition. Finally, note that – for each label – the numbers reported in the first and second row immediately below the label add up to 1.

choices, across our two country samples and three information conditions. Specifically, note that the left panel classifies strategy options based on the subjects' vocabularies (i.e., *American-* and *British-*English respectively for US and UK participants, as shown in Table 3), regardless of their information condition. Thus, the table gives us an idea of the overall impact of word frequency: as can be seen in the left panel, a plurality of the choices (41.25%) consist of the option with the highest NGRAM value in the responding subject's vocabulary. A Hotelling's T-squared generalized means test (conducted on the sample of per-subject mean choices) confirms that the distribution of choices differs from the fully mixed equilibrium assigning equal probability to all the strategies ($N = 80$ obs., $T^2 = 15.14$, $p = 0.001$). Moreover, a two-tailed Wilcoxon signed-rank test indicates that the frequency of choice of the option with the highest NGRAM value is significantly different from chance ($N = 80$ obs., $z = 3.467$, $p = 0.000$).

For the purpose of comparing choice distributions between countries, the right panel of Table 4 classifies strategy options based on the labels' word frequency in the *American-English* vocabulary only. By taking a look at the right panel (Table 4), the reader will notice that the distribution of choices varies by country. In particular, the higher the n-gram frequency of a label in the American-English vocabulary, the more likely it is for the associated option to be selected by participants in the US rather than in the UK (remarkably, the option with the highest American-English NGRAM was chosen 47.25% and 38.75% of the time by US and UK residents, respectively). A two-group Hotelling's T-squared generalized means test (conducted on the samples of per-subject mean choices) confirms that the US and UK choice distributions differ from each other ($N = 80$ obs., $T^2 = 11.665$, $p = 0.004$). (The econometric analysis below will delve into the British-English vocabulary as well.)

To sum up, the above provides evidence in support of H4: behavior differs between US and UK residents, and overall is positively related to the labels' n-gram frequency in the vocabulary of the respective countries. (For a more granular breakdown of the data with respect to both the *American-* and *British-*English vocabularies, see Table 3A in Appendix A.) That said, we note that since the above tests are conducted on the sample of mean choices, they do not account for differences in individual responses across games; also, the above tests do not control for differences in the labels' position on the screen or their numerical NGRAM values. For these reasons, later on we will corroborate our findings by discussing some robustness checks.

Choice by word frequency	Pooled country samples	US residents	UK residents
	Labels' word frequency refers to the relevant vocabulary (i.e., <i>American-</i> and <i>British-</i> English for US and UK residents, respectively)	Labels' word frequency refers to the <i>American-</i> English vocabulary	
Strategy-label with <i>highest</i> NGRAM metric is chosen [f_H], %	41.25 (.1951)	47.25 (.1986)	38.75 (.1785)
Strategy-label with <i>middling</i> NGRAM metric is chosen [f_M], %	28.13 (.1599)	30.25 (.1656)	26.00 (.1532)
Strategy-label with <i>lowest</i> NGRAM metric is chosen [f_L], %	30.62 (.1871)	22.50 (.1597)	35.25 (.1739)
Total, %	100	100	100
Total # triplets	800	400	400
Subjects	80	40	40

Table 4 - (Per-subject) mean choices, given a classification of the labels based on their n-gram frequency; in parentheses is the standard deviation. Note: the number of triplets equals the number of games (i.e., 10) times the number of participants in each condition. The *left panel* presents summary data by pooling choices from the US and UK samples (in each country sample, strategy options are ranked by word frequency in the respective vocabulary). For the sole purpose of comparing distributions across samples, the *right panel* breaks down the data by country, given a classification of the labels based on the American-English corpus.

We move on to address H5, which concerns differences across information conditions (as opposed to differences across country samples). To that end, Table 5 below divides the data into three mutually exclusive groups. The first group (“know-SAME”) consists of participants who were informed that their assigned partner resided in the same country (i.e., US participants in the *know-US* condition, and UK participants in the *know-UK* condition). The second group consists of subjects who were not informed about their partner’s residence, and corresponds to all the participants in the *No-info* condition. The third group (“know-OTHER”) consists of participants who were informed that their assigned partner resided in a different country (i.e., US participants in the *know-UK* condition, and UK participants in the *know-US* condition).

	<i>Know-SAME</i>	<i>NO-Info</i>	<i>Know-OTHER</i>
Choice by word frequency			
Strategy-label with <i>highest</i> NGRAM metric is chosen [f_H], %	50.00 (.1809)	41.62 (.1818)	30.50 (.1904)
Strategy-label with <i>middling</i> NGRAM metric is chosen [f_H], %	29.13 (.1621)	27.29 (.1627)	28.50 (.1598)
Strategy-label with <i>lowest</i> NGRAM metric is chosen [f_H], %	20.87 (.1311)	31.09 (.1882)	41.00 (.1889)
Total, %	100	100	100
Total # triplets	230	370	200
Subjects	23	37	20

Table 5 - (Per-subject) mean choices, given a classification of the labels based on the relevant n-gram frequency (i.e., with respect to the American- and British-English vocabularies for US and UK residents, respectively, regardless of their information condition). In parentheses is the standard deviation. Note: “Know-SAME” includes US residents in the *know-US* condition, and UK residents in the *know-UK* condition. Participants in “NO-Info” received no information about the partner’s country of residence. “Know-OTHER” includes US residents in the *know-UK* condition, and UK residents in the *know-US* condition.

As can be seen in Table 5 above, the mean distribution of choices varies with each group. A quick glance reveals that the *most* frequently-mentioned label (i.e., the option with the highest NGRAM in the responding subject’s vocabulary) was chosen less and less often when moving from know-SAME (50.00%) to No-info (41.62%) to know-OTHER (30.50%). A Kruskal-Wallis test (conducted on the sample of per-subject mean choices) confirms significant differences in the choice of the option with the highest NGRAM across groups ($N = 80$ obs., $\chi^2_2 = 10.597$, $p = 0.005$, two-tailed). Further, Table 5 reveals that the *least* frequently-mentioned label (i.e., the option with the lowest NGRAM) was chosen more and more often, moving from know-SAME (20.87%) to No-info (31.09%) to know-OTHER (41.00%), with such differences being again significant ($N = 80$ obs., $\chi^2_2 = 12.032$, $p = 0.002$, two-tailed Kruskal-Wallis test).²⁷ The above provides some preliminary evidence in support of H5.

²⁷ Also, the distribution of choices in each of the know-SAME and No-info groups differs from the fully mixed equilibrium assigning equal probability to all the strategies, as is confirmed by Hotelling’s T-squared generalized means tests conducted on the samples of mean choices (for *know-SAME*: $N = 23$ obs., $T^2 = 26.91$, $p = 0.000$; for

The tests presented so far provide some insights into the behavioral patterns that emerge from a dataset consisting of average choices (i.e., to satisfy the assumption of independence of observations, the tests above are conducted on the sample of per-subject mean choices, as described in footnote 15, p. 11). In order to account for any differences in choice behavior across games, we proceed to corroborate our findings by conducting an econometric analysis of the full sample of individual observations (while adjusting standard errors for clustering on the subjects). We start by investigating the relationship between one’s use of word frequency and one’s knowledge of the counterpart’s country of residence. To that end, we coded a “knowledge” ordinal variable as follows: this variable takes on value 1 if a subject is in the *know-OTHER* group (i.e., a subject knows that the partner resides in a different country); it takes on value 2 if a subject is in the *NO-info* condition (i.e., a subject receives no information about the partner’s country of residence); it takes on value 3 if a subject is in the *know-SAME* group (i.e., a subject knows that the partner resides in the same country). We then coded a binary “NGRAM prediction” variable, which takes on value 1 if a subject selects the option with the highest NGRAM value in her own vocabulary, and takes on value 0 otherwise. A logit model consisting of the “NGRAM prediction” as the binary dependent variable (and of the “knowledge” variable as the sole predictor) confirms a significant positive effect of the knowledge variable on the choice of the option with the highest NGRAM (for the *US sample*, $\text{coef.} = .4529, z = 2.65, p = 0.008$, two-tailed logit with standard errors clustered on 40 subjects; for the *UK sample*, $\text{coef.} = .3851, z = 2.44, p = 0.015$, two-tailed logit with standard errors clustered on 40 subjects). The above confirms that – in each of our two country samples – *subjects used word frequency in a strategic manner*. The more they had reason to believe that their partner was alike, the more likely they were to select the most frequently-mentioned label in their respective vocabularies. So, the results strengthen our previous evidence in support of H5.

We turn to the next robustness checks. (We shall report the main findings here, while we refer the reader to Appendix A for the econometric tables and further commentary.) In order to control for differences in the labels’ position or differences in their numerical NGRAM values, we ran alternative-specific conditional logit models (“asclogit”; McFadden’s choice model, 1973). Overall, the models confirm a significant positive effect of word frequency on choice

No-info: $N = 37$ obs., $T^2 = 9.21, p = 0.018$). However, we find no such difference in *know-OTHER* ($N = 20$ obs., $T^2 = 3.73, p = 0.199$). The latter suggests that subjects picked an option at random only if they did not have a compelling reason to rely on the labels’ word frequency in their own vocabularies.

behavior. In particular, the higher the n-gram frequency of a label in the country's relevant vocabulary, the more likely it is for that option to be selected (regardless of the label's position on the screen). Moreover, when comparing choice behavior in the *know-SAME* and *know-OTHER* groups, the models reveal significant differences in the relative impact of word frequency; that is, if one is informed that the partner resides in a different country, then one is less likely to select the option with the highest NGRAM in one's vocabulary (see models [5]-[6] of Table 4A, in Appendix A).

We conclude this section by discussing coordination rates. For this purpose, we compared the (expected) coordination rates that would be obtained if participants in different subsamples were paired with each other, using Monte Carlo methods (see footnote 17, p. 11). In short, whereas picking at random implies a coordination rate of 0.33, our subjects' choice behavior implies an expected coordination rate of 0.55 and 0.43 respectively for US and UK participants *knowingly* paired with their compatriots. Notably, these rates drop to 0.47 (0.39) in the case of US (UK) participants *unknowingly* paired with their compatriots. We finally considered the case in which US and UK participants were *knowingly* and *unknowingly* paired with each other, and found that the expected coordination rates were respectively 0.37 and 0.38. In summary, our results show that successful coordination is more likely when subjects are *knowingly* paired with partners from their own country, as opposed to when they are *knowingly* or *unknowingly* paired with partners from a different country.

VI - Conclusion

We have presented a set of studies that test whether the frequency with which labels are mentioned in everyday language may affect game play. In the first study, we found that the labels' frequency of occurrence in the vocabulary of the subject (quantified by the Google Books n-gram frequency) is a good predictor of choice behavior in coordination games. Our second study verifies if subjects utilize word frequency strategically rather than naïvely. To do so, we contrasted participants in coordination games with participants in three alternative roles, namely, "Pickers", "Hiders", and "Seekers". The data reveal that Pickers are as likely as Coordinators to select the most frequently-mentioned label; instead, Hiders are less likely than Seekers and, in turn, Seekers are less likely than Coordinators. This pattern suggests a (boundedly) rational use of frequently-mentioned labels. Our third study delves into the strategic use of word frequency in coordination games, by contrasting culturally diverse participants. To that end, we recruited

samples of US and UK participants, and then varied their knowledge of the counterpart's country of residence. Consistent with our predictions, we found that behavior differs across US and UK residents, and it is positively related to the labels' n-gram frequency in the subject's own vocabulary. Further, a subject is less likely to rely on word frequency as a means to guiding her behavior *if* she knows that the counterpart resides in a different country.

Our approach was inspired by previous evidence on the use of word frequency as a cue in *non*-strategic tasks (e.g., Goldstein and Gigerenzer, 2002; Dougherty et al., 2008). Here, we have shown that word frequency has a role in strategic reasoning too. Remarkably, our results imply that subjects are consciously aware as to how labels might be perceived by culturally alike counterparts, and accordingly adjust their strategies. In this regard, our results are related to a stream of research suggesting that individuals are aware of the communication difficulties that arise when groups with different “conversational codes” merge with each other (Weber and Camerer, 2003; Feiler and Camerer, 2010). It is also worth noting a connection with *rational speech act theory*, which formalizes how participants in conversational interactions make inferences about the meaning of utterances so as to achieve “coordination of meaning”, based on their knowledge of the counterpart and context (Goodman and Frank, 2016).

Before concluding, we note that the past decade has seen the growth of large-scale online datasets and, with it, unique opportunities to investigate human behavior and its cultural correlates. In particular, the Google Books corpus has been used to analyze trends in stereotypes and wellbeing across cultures (e.g., Garg et al., 2018; Hills et al., 2019). More generally, internet data on human activity have been used in fields as diverse as public health (Hawn, 2009), cognitive science (Griffiths, 2015), and management (George, Osinga, Lavie, and Scott, 2016). Our paper shows that such data may inform theories of strategic reasoning as well. Relatedly, we note that although this paper has focused on strategic problems with incentives *to* and *not to* coordinate, our characterization of prominence may have wider application; in fact, subjects' exposure to alternative labels – *ceteris paribus* – can inform behavioral predictions in any class of games with strategic uncertainty.

To conclude, this paper has proposed and tested an a-priori measurable proxy for prominence that explicitly rests on players' culture. The results provide very robust evidence in support of our characterization of prominence. In doing so, the results contribute to shedding light on the relationship between culture, (bounded) rationality and coordination, which plays an important role in several interactions among economic agents.

APPENDIX A

Additional figures and tables

Study 1

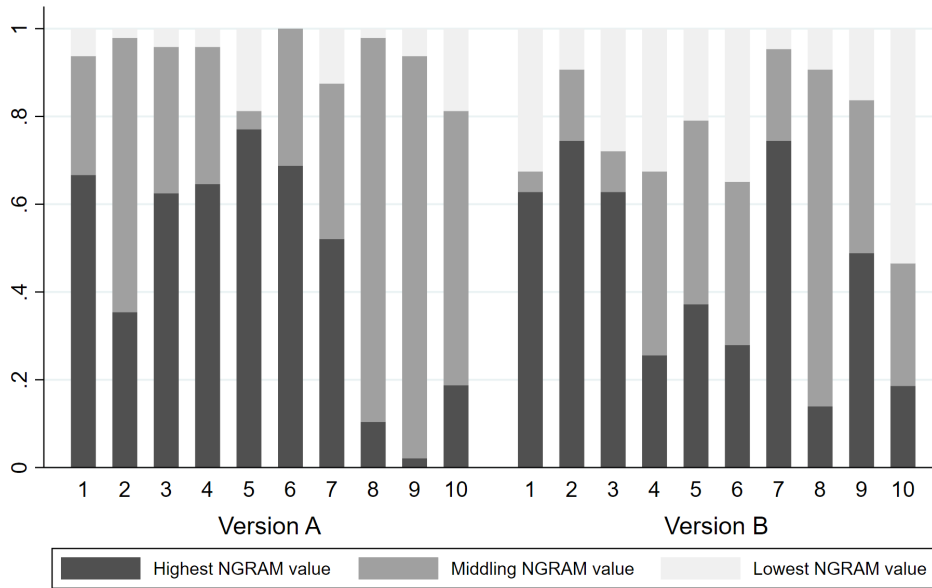


Figure 1A - Frequency distributions of individual-level choices in each of the 1-10 games in Version A (left panel) and Version B (right panel). For visual clarity, in each of the bars, the options have been arranged according to their NGRAM values: darker shades of gray represent options with relatively higher NGRAM values. For the list of labels in each of the games, see Table 1 in the body of the paper.

Study 2




Figure 2A - Frequency distributions of individual-level choices in each of the *1-10 games* and for each of the *1-4 conditions*. Conditions 1-4 refer to “Coordinate”, “Pick”, “Seek”, and “Hide”, respectively. For visual clarity, in each of the bars, the options have been arranged according to their NGRAM values (i.e., darker shades of gray represent options with relatively higher NGRAM values). For the list of labels in each of the 1-10 games, see the left panel of Table 1 in the body of the paper.

We now present a set of alternative-specific conditional logit models (“asclogit”; McFadden’s choice model, 1973). Before describing the list of predictors, note that asclogit analysis differs from other logistic models in that it accounts for “alternative-specific” variables, as well as “case-specific” variables, as defined below.

Alternative-specific variables represent attributes that may vary across each of the options in a choice task (e.g., the labels’ numerical NGRAM value varies across the three options). If an alternative-specific predictor has a positive coefficient, then it implies that the alternative featuring a higher value is more likely selected than the other alternatives, whichever they are. Relatedly, note that alternative-specific predictors do not have an explicit base category; for further discussion, see Cameron and Trivedi (2010, pp. 503-511).

Case-specific variables, on the other hand, represent attributes that are common to each of the three options in a given choice task (e.g., the experimental condition). Coefficients of case-specific predictors are interpreted as parameters of an ordinary multinomial logit model against the base category. Note that the following tables present an upper section, which lists alternative-specific predictors, and a lower bifurcated section that lists case-specific predictors.

In particular, our models consist of the following predictors: *i.* a continuous variable “WF” representing the word frequency of a label (quantified by the NGRAM values in Table 1 in the body of the paper); *ii.* a dummy variable “C” indicating the experimental condition to which a triplet belongs, as specified in the table notes below; *iii.* an interaction variable “WFC” (which captures how the impact of a change in a label’s NGRAM varies with the experimental condition).

	[1] <i>Pick & Coordinate</i>		[2] <i>Coordinate & Seek</i>	
Choice of option				
Relative word frequency (“WF”)	.637*** (.161)		.707*** (.111)	
Relative word frequency * <i>Condition</i> (“WFC”)	-.050 (.104)		-.119** (.049)	
	Center label	Bottom label	Center label	Bottom label
Condition (“C”)	-.948*** (.212)	-.853*** (.212)	1.157*** (.227)	.896*** (.214)
Constant	-.357*** (.115)	-.098 (.110)	-1.305*** (.177)	-.951*** (.181)
Log pseudolikelihood	-746.093		-824.402	
Obs.	810		860	

Table 1A - Alternative-specific conditional logit model coefficients; in parentheses are standard errors clustered on the subjects (*, **, and *** indicate $p < 0.10$, $p < 0.05$ and $p < 0.01$, respectively, for the relevant Z Statistic, two-tailed tests).

Model [1]: *C* takes on value 0 if a subject is assigned to Pick, and value 1 if assigned to Coordinate.

Model [2]: *C* takes on value 0 if a subject is assigned to Coordinate, and value 1 if assigned to Seek.

We begin by discussing model [1] – in the left panel of Table 1A – which uses the full set of option triplets (i.e., ten games per subject) from the Pick and Coordinate conditions, with standard errors clustered on the subjects. First of all, *WF* confirms a significant positive impact of word frequency on choice behavior. This means that the higher the NGRAM value of a label, the more likely it is for that option to be selected, regardless of other attributes such as the experimental condition or the labels’ position (i.e., top, center, or bottom of the list). We further note that *WFC* is non-significant in model [1]. This means that there is no difference in the relative impact of word frequency between Pickers and Coordinators.

A few more comments are in order. The coefficients for dummy C (i.e., $-.948$ and $-.853$) of model [1] capture the impact of a change in experimental conditions on the attractiveness of the center and bottom options (against the top option): they are both negative, indicating that center and bottom options are selected less often by Coordinators than Pickers. Finally, the two *constant* terms (i.e., $-.357$ and $-.098$) of model [1] respectively capture the attractiveness of the center and bottom options (against the top option), due to unmeasured attributes of the alternatives. Both terms are negative (although one is non-significant), reflecting the greater attractiveness of the options located at the top of the list.

Next, model [2] – in the right panel of Table 1A – uses the full set of triplets from the Coordinate and Seek conditions. Model [3] in Table 2A below (left panel) uses the full set of triplets from the Pick and Hide conditions, whereas model [4] (right panel) uses the full set of triplets from the Seek and Hide conditions. Note that WF is positive and significant in all such models. Also note that WFC is negative and significant in models [2], [3], [4]: this implies that the most frequently-mentioned label is selected less often when moving from Coordinate to Seek (model [2]), from Pick to Hide (model [3]), and from Seek to Hide (model [4]), thereby confirming our stated hypotheses.

	[3] <i>Pick & Hide</i>		[4] <i>Seek & Hide</i>	
Choice of option				
Relative word frequency ("WF")	.940*** (.171)		.991*** (.350)	
Relative word frequency * <i>Condition</i> ("WFC")	-.201*** (.053)		-.214** (.100)	
	Center label	Bottom label	Center label	Bottom label
Condition ("C")	.535*** (.174)	.169 (.168)	.325* (.192)	.126 (.172)
Constant	-.357*** (.115)	-.098 (.110)	-.148 (.141)	-.054 (.115)
Log pseudolikelihood	-767.819		-846.129	
Obs.	740		790	

Table 2A - Alternative-specific conditional logit model coefficients; in parentheses are standard errors clustered on the subjects (*, **, and *** indicate $p < 0.10$, $p < 0.05$ and $p < 0.01$, respectively, for the relevant Z Statistic, two-tailed tests).

Model [3]: C takes on value 0 if a subject is assigned to Pick, and value 1 if assigned to Hide.

Model [4]: C takes on value 0 if a subject is assigned to Seek, and value 1 if assigned to Hide.

Study 3




Figure 3A - Frequency distributions of individual-level choices for each of the *1-10 games* and for each of the *1-3 information conditions*. Conditions 1-3 refer to “No-info”, “know-UK”, and “know-US”, respectively. In each of the bars, the options have been arranged according to their NGRAM values (i.e., darker shades of gray represent options with relatively higher NGRAM values in the *relevant* vocabulary; that is, American- and British-English for US and UK residents, respectively). For the list of labels in each of the 1-10 games, see Table 3 in the body of the paper.

Choice by word frequency	US residents			UK residents		
	<i>NO-Info</i>	<i>Know-UK</i>	<i>Know-US</i>	<i>NO-Info</i>	<i>Know-UK</i>	<i>Know-US</i>
Strategy-label with <i>highest</i> NGRAM metric is chosen [f_H], %	47.89 (.1812)	35.00 (.2173)	57.27 (.1618)	35.00 (.1617)	43.33 (.1775)	26.00 (.1577)
Strategy-label with <i>middling</i> NGRAM metric is chosen [f_M], %	31.05 (.1663)	32.00 (.1475)	27.27 (.1902)	23.33 (.1533)	30.83 (.1378)	25.00 (.1715)
Strategy-label with <i>lowest</i> NGRAM metric is chosen [f_L], %	21.06 (.1559)	33.00 (.1828)	15.46 (.0934)	41.67 (.1617)	25.84 (.1443)	49.00 (.1663)
Total, %	100	100	100	100	100	100
Total # triplets	190	100	110	180	120	100
Subjects	19	10	11	18	12	10

Table 3A - (Per-subject) mean choices, given a classification of the labels based on the *relevant* n-gram frequency; that is, American- and British-English vocabularies for US and UK residents, respectively (regardless of their information condition). In parentheses is the standard deviation. For a description of the experimental conditions, see p. 21 above.

For robustness purposes, Table 4A below presents alternative-specific conditional logit models (McFadden, 1973; for general commentary on asclogit analysis, see p. 30 above). In short, model [5] involves US residents' choices, whereas [6] involves UK residents' choices. Each model uses observations from the *know-UK* and *know-US* conditions (as usual, standard errors are clustered on the subjects). Both models consist of the following predictors: *i.* a continuous variable "WF" representing the word frequency of a label in the subject's own vocabulary (quantified by the relevant NGRAM values in Table 3 in the body of the paper); *ii.* a dummy "C" taking on value 0 (1) if a subject is informed that the counterpart resides in the same (different) country; *iii.* an interaction variable "WFC".

A glance at Table 4A reveals that *WF* is positive and significant in both models: this means that the higher the NGRAM value of a label (in the subject's own vocabulary), the more likely it is for that option to be selected. Also, note that *WFC* is negative and significant in both models: this implies that if one is informed that the counterpart resides in a different country, one will less likely select the option with the highest NGRAM in one's vocabulary.

Choice of option	[5] <i>US residents</i>		[6] <i>UK residents</i>	
	Center label	Bottom label	Center label	Bottom label
Relative word frequency ("WF")	4.479*** (.747)		1.841** (.861)	
Relative word frequency * <i>Condition</i> ("WFC")	-4.316*** (1.471)		-4.218*** (1.436)	
Condition ("C")	.193 (.550)	.478 (.508)	-.261 (.412)	-.750* (.437)
Constant	-.491 (.395)	-1.078*** (.316)	-.233 (.313)	-.496 (.329)
Log pseudolikelihood	-203.014		-220.462	
Obs.	210		220	

Table 4A - Alternative-specific conditional logit model coefficients; in parentheses are standard errors clustered on the subjects (*, **, and *** indicate $p < 0.10$, $p < 0.05$ and $p < 0.01$, respectively, for the relevant Z Statistic, two-tailed tests).

Model [5]: US sample; *WF* refers to the American-English vocabulary. Dummy *C* takes on value 0 if a subject is assigned to *know-US*, and value 1 if assigned to *know-UK*.

Model [6]: UK sample; *WF* refers to the British-English vocabulary. Dummy *C* takes on value 0 if a subject is assigned to *know-UK*, and value 1 if assigned to *know-US*.

APPENDIX B

Experimental instructions and screen shots

Study 1

INSTRUCTIONS

Welcome. In this study we are interested in how people coordinate. Please read the following instructions carefully. Note that you may earn a bonus payment in the amount of £0.10 per question, as follows. Specifically, the bonus payment will be determined by your own choices and those made by some other participant, according to the rules described below.

In this experiment each participant (a *Prolific* worker) will be given the same choice tasks, and each participant will be assigned a partner at random (another *Prolific* worker). Note that you will not know the identity of your partner nor will you be able to communicate with him/her.

Both you and your partner will be given 10 multiple-choice questions, with three options in each. If you both select the same option in a given question you will have "coordinated." If you select different options in a given question you will have "miscoordinated." Your goal is to select the same option as your partner (that is, to coordinate with your partner) on as many questions as possible. Recall that you will not know the identity of your partner nor will you be able to communicate with him/her.

Hence, consider a question that asks you to select an option among three, for example, *green*, *blue* and *yellow*. If both you and your partner select the same option (i.e. you both select green, or you both select blue, or you both select yellow) then you will have coordinated. If you select different options (e.g. you select green and your partner selects blue) then you will have miscoordinated.

In short, you will receive a bonus payment of £0.10 per coordinated answer. So - given that this study comprises 10 questions - the total range of bonus payments you can expect to receive is [£0, £1].

>>

Below is an example of an experimental task.

Please choose one option. Each of you and your partner will receive £0.10 if you both choose the same option, £0 otherwise.

- butter
- peppermint
- cocoa

>>

Study 2

Coordinate

Same instructions as in Study 1 above.

Pick

INSTRUCTIONS

Welcome. In this study we are interested in how people pick items. You will be given 10 multiple-choice questions, with three options in each. Your goal is to select one option.

[>>](#)

Below is an example of an experimental task in the *Pick* condition.

Please pick one option.

Nepal

Uzbekistan

Thailand

[>>](#)

Hide-and-Seek

INSTRUCTIONS

Welcome. In this study we are interested in how people choose among options. Please read the following instructions carefully. Note that you may earn a bonus payment in the amount of £0.10 per question, as follows. Specifically, the bonus payment will be determined by your own choices and those made by some other participant, according to the rules described below.

In this experiment each participant (a *Prolific* worker) will be randomly assigned one of two roles (either "Hider" or "Seeker") and matched with a partner of a different role (another *Prolific* worker). Note that you will not know the identity of your partner nor will you be able to communicate with him/her.

Both you and your partner will be given 10 multiple-choice questions, with three options in each. In each question (whether you are assigned the role of Hider or Seeker) you will have to select one of the three options.

Specifically, if you are assigned the role of *Seeker*, your objective will be to select the *same* option as your partner.

Instead, if you are assigned the role of *Hider*, your objective will be to select a *different* option from your partner.

(Recall that you will not be able to communicate with him/her.)

For example, consider a question that asks you to select an option among, say, *green*, *blue* and *yellow*. If both you and your partner select the same option (i.e. you both select green, or you both select blue, or you both select yellow), then the Seeker will receive a bonus payment of £0.10 while the Hider will receive nothing.

On the other hand, if you and your partner select a different option (e.g. you select green and instead your partner selects blue or yellow), then the Hider will receive a bonus payment of £0.10 while the Seeker will receive nothing.

In summary, you may receive a bonus payment of £0.10 per answer depending on your role and the choices of you and your partner. Hence - given that this study comprises 10 questions - the total range of bonus payments you can expect to receive is [£0, £1].

>>

Below is an example of an experimental task in the *Seek* condition.

You have been assigned the role of SEEKER.

Please choose one option. You will receive £0.10 if you choose the *same* option as your partner, £0 otherwise.

Nepal

Uzbekistan

Thailand

Below is an example of an experimental task in the *Hide* condition.

You have been assigned the role of HIDER.

Please choose one option. You will receive £0.10 if you choose a *different* option from your partner, £0 otherwise.

Nepal

Uzbekistan

Thailand

Study 3

INSTRUCTIONS

Welcome. In this study we are interested in how people coordinate. Please read the following instructions carefully. Note that you may earn a bonus payment in the amount of £0.10 per question, as follows. Specifically, the bonus payment will be determined by your own choices and those made by some other participant, according to the rules described below.

In this experiment each participant (a *Prolific* worker) will be given the same choice tasks, and each participant will be assigned a partner at random (another *Prolific* worker). Note that you will not know the identity of your partner nor will you be able to communicate with him/her.

Both you and your partner will be given 10 multiple-choice questions, with three options in each. If you both select the same option in a given question you will have "coordinated." If you select different options in a given question you will have "miscoordinated." Your goal is to select the same option as your partner (that is, to coordinate with your partner) on as many questions as possible. Recall that you will not know the identity of your partner nor will you be able to communicate with him/her.

Hence, consider a question that asks you to select an option among three, for example, *green*, *blue* and *yellow*. If both you and your partner select the same option (i.e. you both select green, or you both select blue, or you both select yellow) then you will have coordinated. If you select different options (e.g. you select green and your partner selects blue) then you will have miscoordinated.

In short, you will receive a bonus payment of £0.10 per coordinated answer. So - given that this study comprises 10 questions - the total range of bonus payments you can expect to receive is [£0, £1].

>>

Below is an example of an experimental task in the *No-info* condition.

Please choose one option. Each of you and your partner receive £0.10 if you both choose the same option, £0 otherwise.

- peach
- pineapple
- pear

>>

Below is an example of an experimental task in the *know-UK* condition.

Please choose one option. Each of you and your partner receive £0.10 if you both choose the same option, £0 otherwise.

Your partner is a Prolific worker who resides in the UK. Your partner may or may not know where you reside.

peach

pineapple

pear

[>>](#)

Below is an example of an experimental task in the *know-US* condition.

Please choose one option. Each of you and your partner receive £0.10 if you both choose the same option, £0 otherwise.

Your partner is a Prolific worker who resides in the US. Your partner may or may not know where you reside.

peach

pineapple

pear

[>>](#)

References

- Alesina, Alberto, Arnaud Devleeschauwer, William Easterly, Sergio Kurlat and Romain Wacziarg.** 2003. "Fractionalization" *Journal of Economic Growth*, 8(2): 155-194.
- Alesina, Alberto and Paola Giuliano.** 2015. "Culture and Institutions" *Journal of Economic Literature*, 53(4): 898-944.
- Alós-Ferrer, Carlos and Christoph Kuzmics.** 2013. "Hidden Symmetries and Focal Points" *Journal of Economic Theory*, 148(1): 226-258.
- Anderson, John R. and Lael J. Schooler.** 1991. "Reflections of the Environment in Memory" *Psychological Science*, 2(6): 396-408.
- Arad, Ayala and Ariel Rubinstein.** 2012. "The 11-20 Money Request Game: A Level-k Reasoning Study" *American Economic Review*, 102(7): 3561-73.
- Bacharach, Michael.** 1993. "Variable Universe Games" in *Frontiers of Game Theory*, ed. K. Binmore et al. Cambridge, MA: MIT Press.
- Bacharach, Michael.** 2006. *Beyond Individual Choice: Teams and Frames in Game Theory*, ed. N. Gold and R. Sugden. Princeton University Press.
- Bacharach, Michael and Michele Bernasconi.** 1997. "The Variable Frame Theory of Focal Points: An Experimental Study" *Games and Economic Behavior*, 19(1): 1-45.
- Bacharach, Michael and Dale O. Stahl.** 2000. "Variable-Frame Level-n Theory" *Games and Economic Behavior*, 32(2): 220-246.
- Balota, David A. and Daniel H. Spieler.** 1999. "Word Frequency, Repetition, and Lexicality Effects in Word Recognition Tasks: Beyond Measures of Central Tendency" *Journal of Experimental Psychology: General*, 128(1): 32-55.
- Bardsley, Nicholas, Judith Mehta, Chris Starmer and Robert Sugden.** 2010. "Explaining Focal Points: Cognitive Hierarchy Theory versus Team Reasoning" *Economic Journal*, 120(543): 40-79.
- Bardsley, Nicholas and Aljaž Ule.** 2017. "Focal Points Revisited: Team Reasoning, the Principle of Insufficient Reason and Cognitive Hierarchy Theory" *Journal of Economic Behavior & Organization*, 133: 74-86.
- Blume, Andreas.** 2000. "Coordination and Learning with a Partial Language" *Journal of Economic Theory*, 95(1): 1-36.
- Camerer, Colin F., Teck-Hua Ho and Juin-Kuan Chong.** 2004. "A Cognitive Hierarchy Model of Games" *Quarterly Journal of Economics*, 119(3): 861-898.
- Cameron, A. Colin and Pravin K. Trivedi.** 2010. *Microeconometrics Using Stata, Revised Edition*. College Station, TX: Stata Press.
- Casajus, André.** 2000. "Focal Points in Framed Strategic Forms" *Games and Economic Behavior*, 32(2): 263-291.
- Charness, Gary, Uri Gneezy and Michael A. Kuhn.** 2013. "Experimental Methods: Extra-Laboratory Experiments—Extending the Reach of Experimental Economics" *Journal of Economic Behavior & Organization*, 91: 93-100.
- Charness, Gary and Alessandro Souto.** 2019. "The Doors of Perception" *ESI Working Paper* 19-32.
- Cooper, Russell W., Douglas V. DeJong, Robert Forsythe and Thomas W. Ross.** 1990. "Selection Criteria in Coordination Games: Some Experimental Results" *American Economic Review*, 80(1): 218-233.
- Costa-Gomes, Miguel A. and Vincent P. Crawford.** 2006. "Cognition and Behavior in Two-Person Guessing Games: An Experimental Study" *American Economic Review*, 96(5): 1737-1768.
- Costa-Gomes, Miguel A., Vincent P. Crawford and Bruno Broseta.** 2001. "Cognition and Behavior in Normal-Form Games: An Experimental Study" *Econometrica*, 69(5): 1193-1235.
- Crawford, Vincent P., Miguel A. Costa-Gomes and Nagore Iriberry.** 2013. "Structural Models of Nonequilibrium Strategic Thinking: Theory, Evidence, and Applications" *Journal of Economic Literature*, 51(1): 5-62.
- Crawford, Vincent P., Uri Gneezy and Yuval Rottenstreich.** 2008. "The Power of Focal Points Is Limited: Even Minute Payoff Asymmetry May Yield Large Coordination Failures" *American Economic Review*, 98(4): 1443-1458.
- Crawford, Vincent P. and Hans Haller.** 1990. "Learning How to Cooperate: Optimal Play in Repeated Coordination Games" *Econometrica*, 58(3): 571-595.

- Crawford, Vincent P. and Nagore Iriberry.** 2007. "Fatal Attraction: Salience, Naivete, and Sophistication in Experimental Hide-and-Seek Games" *American Economic Review*, 97(5): 1731-1750.
- Dougherty, Michael R., Ana M. Franco-Watkins and Rick Thomas.** 2008. "Psychological Plausibility of the Theory of Probabilistic Mental Models and the Fast and Frugal Heuristics" *Psychological Review*, 115(1): 199-211.
- Easterly, William and Ross Levine.** 1997. "Africa's Growth Tragedy: Policies and Ethnic Divisions" *Quarterly Journal of Economics*, 112(4): 1203-1250.
- Einav, Liran and Jonathan Levin.** 2014. "Economics in the Age of Big Data" *Science*, 346(6210).
- Faillo, Marco, Alessandra Smerilli and Robert Sugden.** 2017. "Bounded Best-Response and Collective-Optimality Reasoning in Coordination Games" *Journal of Economic Behavior & Organization*, 140: 317-335.
- Feiler, Lauren and Colin F. Camerer.** 2010. "Code Creation in Endogenous Merger Experiments" *Economic Inquiry*, 48(2): 337-352.
- Garg, Nikhil, Londa Schiebinger, Dan Jurafsky and James Zou.** 2018. "Word Embeddings Quantify 100 Years of Gender and Ethnic Stereotypes" *Proceedings of the National Academy of Sciences*, 115(16): E3635-E3644.
- George, Gerard, Ernst C. Osinga, Dovev Lavie and Brent A. Scott.** 2016. "Big Data and Data Science Methods for Management Research" *Academy of Management Journal*, 59(5):1493-1507.
- Gigerenzer, Gerd and Daniel G. Goldstein.** 1996. "Reasoning the Fast and Frugal Way: Models of Bounded Rationality" *Psychological Review*, 103(4): 650-669.
- Ginsburgh, Victor and Shlomo Weber.** 2020. "The Economics of Language" *Journal of Economic Literature*, 58(2): 348-404.
- Goldstein, Daniel G. and Gerd Gigerenzer.** 2002. "Models of Ecological Rationality: The Recognition Heuristic" *Psychological Review*, 109(1): 75-90.
- Goodman, Noah D. and Michael C. Frank.** 2016. "Pragmatic Language Interpretation as Probabilistic Inference" *Trends in Cognitive Sciences*, 20(11): 818-829.
- Griffiths, Thomas L.** 2015. "Manifesto for a New (Computational) Cognitive Revolution" *Cognition*, 135: 21-23.
- Hargreaves Heap, Shaun, David Rojo Arjona and Robert Sugden.** 2014. "How Portable Is Level-0 Behavior? A Test of Level-k Theory in Games with Non-Neutral Frames" *Econometrica*, 82(3): 1133-1151.
- Hargreaves Heap, Shaun, David Rojo Arjona and Robert Sugden.** 2017. "Coordination When There Are Restricted and Unrestricted Options" *Theory and Decision*, 83(1): 107-129.
- Harsanyi, John C. and Reinhard Selten.** 1988. *A General Theory of Equilibrium Selection in Games*. Cambridge, MA: MIT Press.
- Hawn, Carleen.** 2009. "Take Two Aspirin and Tweet Me in the Morning: How Twitter, Facebook, and Other Social Media Are Reshaping Health Care" *Health Affairs*, 28(2): 361-368.
- Hedden, Trey and Jun Zhang.** 2002. "What Do You Think I Think You Think?: Strategic Reasoning in Matrix Games" *Cognition*, 85(1): 1-36.
- Hertwig, Ralph, Stefan M. Herzog, Lael J. Schooler and Torsten Reimer.** 2008. "Fluency Heuristic: A Model of How the Mind Exploits a By-Product of Information Retrieval" *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 34(5): 1191-1206.
- Hertwig, Ralph, Thorsten Pachur and Stephanie Kurzenhäuser.** 2005. "Judgments of Risk Frequencies: Tests of Possible Cognitive Mechanisms" *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 31(4): 621-642.
- Hills, Thomas T., Eugenio Proto, Daniel Sgroi and Chanuki Illushka Seresinha.** 2019. "Historical Analysis of National Subjective Wellbeing Using Millions of Digitized Books" *Nature Human Behaviour*, 3: 1-5.
- Janssen, Maarten C. W.** 2001. "Rationalizing Focal Points" *Theory and Decision*, 50(2): 119-148.
- Jescheniak, Jörg D. and Willem J. M. Levelt.** 1994. "Word Frequency Effects in Speech Production: Retrieval of Syntactic Information and of Phonological Form" *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20(4): 824-843.
- Kapferer, Jean-Noel.** 2012. *The New Strategic Brand Management: Advanced Insights and Strategic Thinking*. London: Kogan Page.
- Kossoudji, Sherrie A.** 1988. "English Language Ability and the Labor Market Opportunities of Hispanic and East Asian Immigrant Men" *Journal of Labor Economics*, 6(2): 205-228.
- Lang, Kevin.** 1986. "A Language Theory of Discrimination" *Quarterly Journal of Economics*, 101(2): 363-382.
- Lazear, Edward P.** 1999. "Culture and Language" *Journal of Political Economy*, 107(6): S95-S126.
- Lewis, David K.** 1969. *Convention: A Philosophical Study*. Cambridge, MA: Harvard U. Press.

- Lichtenstein, Sarah, Paul Slovic, Baruch Fischhoff, Mark Layman and Barbara Combs.** 1978. "Judged Frequency of Lethal Events" *Journal of Experimental Psychology: Human Learning and Memory*, 4(6): 551-578.
- Luce, R. Duncan and Howard Raiffa.** 1957. *Games and Decisions: Introduction and Critical Survey*. John Wiley & Sons.
- Matz, Sandra C. and Oded Netzer.** 2017. "Using Big Data as a Window into Consumers' Psychology" *Current Opinion in Behavioral Sciences*, 18: 7-12.
- Mauersberger, Felix and Rosemarie Nagel.** 2018. "Levels of Reasoning in Keynesian Beauty Contests: A Generative Framework" in *Handbook of Computational Economics, Vol. 4*, ed. C. Hommes and B. LeBaron. Amsterdam: Elsevier.
- McFadden, Daniel.** 1973. "Conditional Logit Analysis of Qualitative Choice Behavior" in *Frontiers in Econometrics*, ed. P. Zarembka. New York: Academic Press.
- Mehta, Judith, Chris Starmer and Robert Sugden.** 1994. "The Nature of Salience: An Experimental Investigation of Pure Coordination Games" *American Economic Review*, 84(3): 658-673.
- Melitz, Jacques.** 2008. "Language and Foreign Trade" *European Economic Review*, 52(4): 667-699.
- Michel, Jean-Baptiste, Yuan Kui Shen, Aviva Presser Aiden, Adrian Veres, Matthew K. Gray, the Google Books Team, Joseph P. Pickett, Dale Hoiberg, Dan Clancy, Peter Norvig, Jon Orwant, Steven Pinker, Martin A. Nowak and Erez Lieberman Aiden.** 2011. "Quantitative Analysis of Culture Using Millions of Digitized Books" *Science*, 331(6014): 176-182.
- Morrison, Catriona M. and Andrew W. Ellis.** 1995. "Roles of Word Frequency and Age of Acquisition in Word Naming and Lexical Decision" *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21(1): 116-133.
- Nagel, Rosemarie.** 1995. "Unraveling in Guessing Games: An Experimental Study" *American Economic Review*, 85(5): 1313-1326.
- Penczynski, Stefan P.** 2016. "Strategic Thinking: The Influence of the Game" *Journal of Economic Behavior & Organization*, 128: 72-84.
- Rubinstein, Ariel and Yuval Salant.** 2016. "'Isn't Everyone like Me?': On the Presence of Self-Similarity in Strategic Interactions" *Judgment and Decision Making*, 11(2): 168-173.
- Rubinstein, Ariel, Amos Tversky and Dana Heller.** 1997. "Naive Strategies in Competitive Games" in *Understanding Strategic Interaction*, ed. W. Albers et al. Berlin: Springer-Verlag.
- Schelling, Thomas C.** 1960. *The Strategy of Conflict*. Cambridge, MA: Harvard U. Press.
- Schwarz, Norbert, Herbert Bless, Fritz Strack, Gisela Klumpp, Helga Rittenauer-Schatka and Annette Simons.** 1991. "Ease of Retrieval as Information: Another Look at the Availability Heuristic" *Journal of Personality and Social Psychology*, 61(2): 195-202.
- Sebanz, Natalie, Günther Knoblich and Wolfgang Prinz.** 2003. "Representing Others' Actions: Just Like One's Own?" *Cognition*, 88(3): B11-B21.
- Seidenberg, Mark S. and James L. McClelland.** 1989. "A Distributed, Developmental Model of Word Recognition and Naming" *Psychological Review*, 96(4): 523-568.
- Stahl, Dale O. and Paul W. Wilson.** 1995. "On Players' Models of Other Players: Theory and Experimental Evidence" *Games and Economic Behavior*, 10(1): 218-254.
- Sugden, Robert.** 1995. "A Theory of Focal Points" *Economic Journal*, 105(430): 533-550.
- Tversky, Amos and Daniel Kahneman.** 1973. "Availability: A Heuristic for Judging Frequency and Probability" *Cognitive Psychology*, 5(2): 207-232.
- Tversky, Amos and Daniel Kahneman.** 1974. "Judgment under Uncertainty: Heuristics and Biases" *Science*, 185(4157): 1124-1131.
- Van Huyck, John B., Raymond C. Battalio and Richard O. Beil.** 1990. "Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure" *American Economic Review*, 80(1): 234-248.
- Varian, Hal R.** 2014. "Big Data: New Tricks for Econometrics" *Journal of Economic Perspectives*, 28(2): 3-28.
- Weber, Roberto A. and Colin F. Camerer.** 2003. "Cultural Conflict and Merger Failure: An Experimental Approach" *Management Science*, 49(4): 400-415.
- Yin, Zhijun, Lina M. Sulieman, and Bradley A. Malin.** 2019. "A Systematic Literature Review of Machine Learning in Online Personal Health Data" *Journal of the American Medical Informatics Association*, 26(6): 561-576.