A case of constructional contamination in English: Modified noun phrases influence adverb placement in the passive

Martin Hilpert
Susanne Flach
Université de Neuchâtel

Abstract

This paper discusses a case of what Pijpops and Van de Velde (2016) call constructional contamination. Specifically, we investigate the influence of English modified noun phrases on variation in adverb placement in the passive. On the basis of data from the COCA, we argue that highly frequent nominal expressions such as *sexually transmitted disease* influence adverb placement in the passive, which offers speakers a choice between adverb-initial order (*The disease was sexually transmitted*) and adverb-final order (*The disease was transmitted sexually*). Our results thus corroborate findings from Dutch corpora (Pijpops and Van de Velde 2016) and suggest that constructional contamination is a phenomenon that can be observed across different languages. We further discuss the role of constructional contamination for analogy and contrast.

1 Introduction

This paper addresses the overall topic of this volume, contrast and analogy, by discussing a phenomenon that has been labeled *constructional contamination* (Pijpops and Van de Velde 2016). The term describes a relation between two constructions of a language, such that usage frequencies of one construction influence, or "contaminate", patterns of variation in another construction. Constructional contamination is thus evidence for associative links between formally different constructions, and it can provide interesting insights into the way speakers draw analogies between constructions that exhibit superficial similarities, even when they differ in their morphosyntactic structure. The phenomenon of constructional contamination is further particularly relevant for theories such as Construction Grammar

(Goldberg 1995, 2006, Hilpert 2014), which posit that knowledge of language is organized in a network-like fashion. Pijpops and Van de Velde (2016, to appear) introduce the phenomenon through analyses of several examples from Dutch, and they encourage the study of cases from other languages. Taking its cues from this challenge, this paper examines a case of constructional contamination in English. Specifically, the focus is on adverb placement variation in the English passive, as illustrated in the examples below.

(1) Garlic is widely believed to help protect against infections.
Garlic is believed widely to help protect against infections.

(2) Stainless steels are extensively used in petrochemical plants.
Stainless steels are used extensively in petrochemical plants.

As the examples show, speakers are free to vary the order of the past participle, such as *believed* or *used*, and its modifying adverb, such as *widely* or *extensively*. While both orders are possible and fully grammatical, specific combinations of adverbs and participles exhibit strong preferences. In the case of *widely* and *believed*, corpus data from the COCA (Davies 2008) reveal a preference for the adverb-initial order *widely believed*. For *extensively* and *used*, the reverse is true. Speakers strongly prefer the adverb-final order *used extensively*. It could be suggested that these preferences reflect lexical idiosyncrasies that do not have any underlying explanation. In fact, Huddleston and Pullum (2002: 576) note that "[o]nly rather broad and approximate flexible generalisations about adjunct placement and sequence can be made. There is a great deal of variation in use, and features of context, style, prosody, and euphony play a role in some decisions." The argument that is made here is that at least some of the variability in adverb placement in the English passive can be explained by constructional contamination. The analysis in this paper will examine the influence of a potentially contaminating construction, namely a noun phrase construction that involves a sequence of an adverb and a participle as a complex modifier of the head noun. As the

examples below illustrate, the noun phrase construction exclusively permits the adverb-initial order.[1]

(3) a widely believed myth
* a believed widely myth
an extensively used technique
* a used extensively technique

The question that will be pursued in this paper is whether usage frequencies in the modified noun phrase construction can partly explain the variation that we see in adverb placement in the passive construction. Put simply, if a string such as *widely believed* occurs frequently as a modifier in a noun phrase, do these occurrences bias speakers towards the adverb-initial order in the passive? In order to answer this question, this paper will draw on data from the COCA (Davies 2008) and examine different frequency measures of modified noun phrases, i.e. the potentially contaminating construction, and of the passive, i.e. the construction that is affected by this contamination.

The remainder of this paper is organized as follows. Section 2 reviews the concept of constructional contamination in more detail and spells out how Pijpops and Van de Velde (2016, to appear) measure it in corpus data. Section 3 discusses how the data for the present study was collected and what methods were used for the analysis. Section 4 presents the results. Section 5 ties the results to a discussion of contrast and analogy in more general terms, and it offers some concluding remarks.

2 Constructional contamination

[1] We thank Dirk Pijpops (p.c.) for pointing out to us that the adverb-final order is in fact found in nominal constructions with postposed modifiers, as in *a myth widely believed*. We analyze these as reduced relative clauses that share structural characteristics of the passive. For example, the postposed construction can include a phrase with *by*, as in *a myth widely believed by laypeople*. The nominal construction, by contrast, does not allow this: **a widely believed by laypeople myth*.

Pijpops and Van de Velde (2016: 543) define constructional contamination in the following way:

> Constructional contamination is the effect whereby a subset of instances of a target construction is (stochastically) affected in its realization by a contaminating construction, because of a coincidental resemblance between the superficial strings of instances of the target construction and a number of instances of the contaminating construction.

They demonstrate the effect of constructional contamination on the basis of a case study that draws on two Dutch constructions in which the same string of words reflects two different morpho-syntactic structures. Their target construction, i.e. the construction that undergoes constructional contamination, is the Dutch partitive genitive, which is shown below (Pijpops and Van de Velde 2016: 544).

(4)	in	begin	van	de	week	iets	verkeerd	gegeten
	[in	beginning of	the	week]$_{PP}$		[something wrong]$_{NP}$		eaten

'I ate something wrong at the start of the week.'

The Dutch partitive genitive displays variation with regard to an inflectional suffix that can appear on the adjective: besides the form *iets verkeerd*, which is shown in (4), speakers may alternatively produce *iets verkeerds*, which adds an *-s* suffix to the expression. Pijpops and Van de Velde (2016: 545) argue that the variation between the suffixed variant and the bare variant is influenced by a contaminating construction that contains a string that superficially resembles the bare variant of the partitive genitive. This construction is shown in (5).

(5)	dat	iets	verkeerd	geïnterpreteerd	wordt?
	that	[something]$_{NP}$	[wrongly]$_{AdvP}$	interpreted	gets

'...that something gets wrongly interpreted?'

Importantly, the contaminating construction itself is not free to vary. It only allows a bare form, there is no suffixed variant. It is hypothesized that frequent exposure to the string *iets*

verkeerd in the contaminating construction leads speakers to adopt a bias for the bare variant in the partitive genitive.

Pijpops and Van de Velde connect their argument to a constructional view of linguistic knowledge, according to which grammar is organized in a network-like fashion (Langacker 1987, Goldberg 1995). This network contains links of different kinds (Hilpert 2014, Hilpert and Diessel 2017). Most commonly discussed are so-called inheritance links, which connect abstract constructions with more specific instantiations of those constructions. For example, the highly abstract verb phrase construction is connected through inheritance links to more specific verbal constructions, such as the transitive construction or the ditransitive construction. Inheritance links are however not the only type of link. Polysemy links connect the formal pole of a construction with several semantic poles. Links are also assumed between constructions that are synonymous, or related through metaphor (Goldberg 1995: 75). Yet another type of link connects constructions with partial formal or semantic overlap. These links are called subpart links (Hilpert 2014: 62). Subpart links are of particular importance for the concept of constructional contamination, since this phenomenon can be observed with constructions that share aspects of their surface form but that do not mutually instantiate each other through inheritance. If it can be shown that constructions with formal overlap of this kind exert a mutual influence, this would motivate a view of the constructional network that assigns great importance to "horizontal" links, that is, links between constructions that do not stand in a relationship of categorical inclusion. Strong effects of constructional contamination would also align with theories of "good-enough" syntactic processing (Ferreira et al. 2002, Ferreira and Patson 2007), which call into question the idea that speakers perform detailed syntactic parses of sentences that they hear. If speakers process incoming language on the basis of stored chunks, this would explain naturally why frequent strings have an influence even across constructions in which those strings do not have the same syntactic structure.

How can it be established whether there are effects of constructional contamination in a given dataset? Pijpops and Van de Velde (2016) consider different ways of quantifying constructional contamination, several of which target the specific characteristics of the Dutch partitive genitive construction. One of their proposed measures does however lend

itself to an application in different contexts. Their measure of *string resemblance* (2016: 564) is computed as a ratio that is based on the frequency of the contaminating construction, divided by a sum that includes the frequency of both variants of the contaminated construction and the frequency of the contaminating construction. A concrete example, based on the English passive sentences that were discussed above, is shown in (6).

(6) $$\frac{\text{frequency of } \textit{widely believed NOUN}}{\text{frequency of passive } \textit{widely believed} + \text{passive } \textit{believed widely} + \textit{widely believed NOUN}}$$

If we take a string such as *widely believed*, this string occurs with a given frequency in nominal constructions such as *a widely believed myth*. The same string also occurs in passive constructions such as *The story was widely believed* or *The story was believed widely*. The division in (6) yields a numerical value between zero and one. Values close to one indicate a high relative frequency of the contaminating noun phrase construction, and hence a high degree of contamination. If string resemblance is high, Pijpops and Van de Velde (2016) argue that speakers are likely to favor the variant of the target construction that is formally identical to the contaminating construction. We will call this variant the isomorphic variant in the following. Conversely, values that are close to zero indicate that the nominal construction is much less frequent than the passive construction, which translates into a very low degree of contamination, and a low likelihood that speakers will produce the "contaminated" variant of the target construction.

In the following section, which discusses our data and methodology, we will come back to this quantitative measure of constructional contamination, and we will compare it to other possible measures.

3 Data and methodology

3.1 Data retrieval

For the present study, the off-line version of the COCA (Davies 2008) was searched for patterns instantiating two different constructions, namely the English passive and a noun phrase construction with a complex modifier. A first set of tag-based searches identified sequences of the passive construction with adverbial modification of the participle. Two variants of the passive construction were retrieved, one in which a form of the copula was followed by an adverb and a participle (e.g. *was dismissed unfairly*), and another one in which a form of the copula was followed by a participle and an adverb (e.g. *was unfairly dismissed*). The search pattern for the second construction identified sequences of an adverb, a participle, and a noun (e.g. *highly trained specialists*). In order to exclude non-target examples, strings were not retrieved if they were preceded by a verb (e.g. *had significantly reduced stress*). Table 1 shows the overall token frequencies for the three search patterns along with the type frequencies of the respective adverbs and participles.

Construction	Tokens	Adverb types	Participle types
adverb-initial passive	314104	2906	5900
adverb-final passive	107211	2450	3645
complex modifier NP	83118	2814	2428

Table 1: Token and type frequencies of the constructions under analysis

The analyses in this paper are based on the combinations of adverb types and participle types and their respective frequencies in the variants of the passive construction and in the complex modifier NP construction. Of particular relevance is the question whether combinations that are frequent in the complex modifier NP construction are also frequent in the adverb-initial passive, but less so in the adverb-final passive. Table 2 offers a view at twenty highly frequent adverb-participle combinations in the database that occur at least once with each constructional variant. The numbers in bold font indicate the construction with the highest token frequency for the respective collocation.

adverb	participle	passive (ADV-PPART)	passive (PPART-ADV)	complex modifier NP
well	known	**1594**	4	110
best	known	**957**	7	212
also	found	**608**	12	169

widely	used	**501**	55	11
randomly	assigned	**444**	33	39
often	called	**407**	1	93
also	included	**283**	2	205
sexually	abused	**391**	5	74
privately	owned	151	6	**290**
publicly	traded	60	2	**380**
well	established	**376**	3	31
highly	regarded	122	2	**265**
dimly	lit	64	1	**297**
randomly	selected	**308**	38	13
clearly	defined	105	6	**241**
democratically	elected	42	12	**284**
hard	hit	91	**245**	1
specifically	designed	**187**	**136**	8
better	prepared	**331**	1	1
often	seen	**310**	3	2

Table 2: Twenty frequent adverb-participle collocations

The table shows that the variants of the passive and the complex modifier NP construction have collocational preferences that engage in different configurations. For example, the collocation *often seen* is highly frequent in adverb-initial passives, but it is near-absent across adverb-final passives and complex modifier NPs. Collocations such as *privately owned* and *highly regarded* are frequent in adverb-initial passives and complex modifier NPs, but they are very infrequent in adverb-final passives. The crucial question with regard to Table 2 is the following: Does a high frequency in the complex modifier NP construction typically imply a frequency imbalance between the two variants of the passive, such that the adverb-initial variant occurs more frequently than would be expected by chance? A look at the numbers in bold font provides some encouragement. Wherever the complex modifier NP construction has the highest token frequency, there is a relative preference for the adverb-initial passive. It is of course necessary to look at all combinations and at the marginal frequencies of the full table to see whether this tendency holds for the dataset as a whole.

3.2 Differences between adverb-initial and adverb-final passive examples

In their study of the Dutch partitive genitive, Pijpops and Van de Velde (2016) start the analysis by contrasting the two variants of the target construction. They do so on the basis of a distinctive collexeme analysis (Gries and Stefanowitsch 2004), which allows them to

demonstrate that there are significant differences in the collocational profiles of the two variants. The same strategy is adopted here. Table 3 shows the twenty combinations of adverbs and participles that are most strongly distinctive for the adverb-initial variant of the English passive construction; Table 4 shows the preferences of the adverb-final variant. We report values of collostructional strength that are based on log likelihood. All combinations are significantly distinctive at the level of $p<0.001$.

combination	observed adverb-initial	expected adverb-initial	observed adverb-final	expected adverb-final	coll.strength (log likelihood)
well known	1594	1191,2	4	406,8	893,59
best known	957	718,6	7	245,4	499,23
also used	800	607,5	15	207,5	361,96
often used	681	521,8	19	178,2	277,89
also found	608	462,2	12	157,8	271,88
also known	507	381,7	5	130,3	255,52
often called	407	304,1	1	103,9	227,97
closely related	449	339,2	6	115,8	216,51
also asked	401	300,4	2	102,6	215,99
significantly related	458	348,9	10	119,1	199,90
well established	376	282,5	3	96,5	194,22
sexually abused	391	295,2	5	100,8	189,86
better prepared	331	247,5	1	84,5	183,69
positively related	367	277,3	5	94,7	176,38
generally considered	312	234,1	2	79,9	164,65
often seen	310	233,3	3	79,7	156,56
better known	273	204,3	1	69,7	149,97
also included	283	212,5	2	72,5	147,98
widely regarded	269	201,3	1	68,7	147,65
also expected	280	210,2	2	71,8	146,26

Table 3: Distinctive collexemes of the adverb-initial variant of the English passive

combination	observed adverb-initial	expected adverb-initial	observed adverb-final	expected adverb-final	coll.strength (log likelihood)
taken seriously	901	229,6	1	672,4	2456,25
used only	471	138	71	404	911,12
treated fairly	205	53,5	5	156,5	516,96
found only	264	80,7	53	236,3	467,81
known only	195	57,3	30	167,7	374,78
used extensively	182	54	30	158	343,02
hit hard	246	85,8	91	251,2	333,84
treated equally	124	31,8	1	93,2	328,38
taken lightly	116	29,8	1	87,2	306,60
based solely	133	36,4	10	106,6	297,46
limited only	146	42,5	21	124,5	285,68
made even	150	45,6	29	133,4	269,06
achieved only	118	33,1	12	96,9	250,01
used primarily	163	55,2	54	161,8	234,42
treated unfairly	111	31,8	14	93,2	224,39
shot twice	85	22,1	2	64,9	214,79
used interchangeably	84	21,9	2	64,1	212,09
made only	145	48,9	47	143,1	210,82
seen only	127	40,2	31	117,8	209,40
taken literally	79	20,4	1	59,6	206,07

Table 4: Distinctive collexemes of the adverb-final variant of the English passive

A comparison of the two tables reveals that the adverb-final variant shows a marked preference for the adverb *only*, which appears seven times, and the participle *treated*, which appears three times. The observation that the two variants of the passive exhibit different collocational preferences justifies a more thorough analysis of constructional contamination as a potential contributor to these differences.

3.3 Data annotation

In order to assess the potential effect of constructional contamination, we combined the examples of both passive variants into a single dataset. We restricted the dataset to combinations that occur in all three structures that we examine: Combinations had to alternate between the two passives in the COCA and they also had to occur with the contaminating complex noun phrase construction. This led to the exclusion of combinations such as *newly elected*, which is not found in the adverb-final passive, *treated differently*, which is not found in the adverb-initial passive, and *further examined*, which does not occur with the complex noun phrase construction. This procedure reduced the dataset to 88341 examples that are distributed over 3785 different combination types. These examples were annotated for the following variables.

3.3.1 Dependent variable: order of adverb and participle

Each example in the dataset was annotated for our dependent variable, that is, the binary distinction between adverb-initial order (*was internationally recognized*) and adverb-final order (*was recognized internationally*).

3.3.2 NP frequency

In addition, we determined for each example in the database how frequently its combination of adverb and participle is attested in the complex NP construction in the COCA. To illustrate, the combination *internationally recognized* is found 159 times in the complex noun phrase construction in the COCA, so each passive example with this pair was annotated with the logged value of that specific frequency.

3.3.3 String resemblance

Following the practice of Pijpops and Van de Velde (2016), we calculated the measure of string resemblance for each adverb-participle combination through a division of NP frequency by the sum of NP frequency, adverb-initial passive frequency, and adverb-final passive frequency. For the combination *internationally recognized*, which occurs 159 times in the NP construction, 63 times in the adverb-initial passive construction, and 23 times in the adverb-final passive construction, the string resemblance value is 159 / (159 + 63 + 23) = 0.649.

3.3.4 Adverb and participle

We further annotated each example in terms of its adverb-participle combination, and individually for its adverb and its participle. This was done with an eye to including adverb and participle as random factors in our regression analyses, which we discuss below.

3.3.5 Mutual association strength between adverb and participle

As another frequency-related variable, we included the mutual association strength between adverb and participle in the contaminating construction as a variable. The inclusion of such a variable can be motivated as follows. A combination such as *randomly selected*, which is not highly frequent in the noun phrase construction (cf. Table 2), may still be strongly entrenched in speakers' minds as a combination in a nominal construction, since the adverb *randomly* is strongly predictive of *selected* as the following item in that construction. In order to control for the possible effect of mutual association strength, we computed a covarying-collexeme analysis (Stefanowitsch & Gries 2005) for our COCA data of the complex noun phrase construction. Table 5 shows the ten combinations with the highest mutual association strength. We report values of collostructional strength that are based on log likelihood. All combinations are significantly distinctive at the level of $p<0.001$.

adverb	participle	adverb frequency	participle frequency	observed	expected	Coll.strength (log likelihood)
publicly	traded	962	495	380	5,7	3028,94
dimly	lit	322	687	297	2,7	2825,82
densely	populated	323	628	285	2,4	2706,54

democratically	elected	321	1053	284	4,1	2337,61
privately	owned	763	839	290	7,7	1773,15
highly	regarded	1751	333	265	7	1753,91
brightly	lit	353	687	211	2,9	1622,80
freshly	squeezed	720	177	168	1,5	1567,68
clearly	defined	541	928	242	6	1506,48
sparsely	populated	229	628	169	1,7	1439,22

Table 5: Top 10 covarying collexemes in the complex NP construction

All examples in our database of passive constructions were annotated with the values of collostructional strength that correspond to the respective combinations of adverb and participle. The values were scaled and centered prior to the statistical modeling that is described below.

3.4 Splitting the data into two datasets

In order to assure the replicability of our analysis, we decided to split our data into two datasets. We randomly assigned each example in the original dataset to one of the two new datasets. This procedure allows us to safeguard against overfitting and to see how robustly the results of one half of the data carry over to the other half. The subsequent analyses are thus based on one dataset with 44168 examples (dataset A) and a second one with 44173 examples (dataset B).

3.5 Data analysis

For the quantitative analysis of our data, we rely on mixed-effects logistic regression (Baayen 2008, Gries 2015) with the aim of predicting the outcome of the dependent variable (order of adverb and participle in the passive) on the basis of two explanatory variables, namely frequency in the noun phrase construction (logged) and string resemblance. These two enter the model as fixed effects. Additionally, we specify random intercepts for adverb and participle types in order to control for item-specific preferences that each adverb or participle may introduce. We use the *lme4* package in R (Bates et al. 2015).

4 Results

A mixed-effects logistic regression finds significant main effects for NP frequency, string resemblance, and collostructional strength. Table 5 summarizes the fixed effects of the model.

	Estimate	Std. Error	z value	Pr(>|z|)
Intercept	1.23	0.08	15.26	< 2e-16 ***
NP frequency (logged)	0.39	0.02	18.49	< 2e-16 ***
string resemblance	-1.93	0.15	-13.29	< 2e-16 ***
collostructional strength	0.14	0.04	3.64	0.00049 ***

Table 5: Fixed effects in the analysis of Dataset A

First, there is an effect of NP frequency in the expected direction. The higher the frequency of an adverb-participle combination in the complex NP construction, the higher the likelihood of a speaker or writer adopting the adverb-initial order in the passive. Second, there is a significant effect of string resemblance. The lower the string resemblance value of a given adverb-participle combination, the higher the likelihood of a passive with adverb-initial order. To illustrate, a combination such as *positively correlated*, which occurs 286 times with the adverb-initial passive and 24 times with the adverb-final passive, but just once with the complex noun phrase construction, has a string resemblance value of 1 / (1 + 286 + 24) = 0.0032. Combinations with this kind of distribution lead to the negative correlation of string resemblance and likelihood of the adverb-initial passive, as we will discuss in more detail in the following section. Third, the model includes a significant effect of collostructional strength. The higher the mutual association between an adverb and a participle, the more likely it is that speakers will use the adverb-initial variant of the passive.

Table 6 shows the classification accuracy that is achieved by the regression model. The model correctly classifies 84.4% of the data, up from a chance baseline of 78.1%, which

represents the ratio of adverb-initial examples in dataset A. As the table shows, most of the model's misclassifications concern adverb-final examples that are wrongly classified as adverb-initial.

		predicted	
		adverb-final	adverb-initial
observed	adverb-final	4496	5169
	adverb-initial	1696	32807

Table 6: Classification accuracy in the analysis of Dataset A

We used the regression model based on Dataset A for the purpose of predicting the dependent variable in Dataset B. Table 7 shows the prediction accuracy for Dataset B, which is only slightly worse than the accuracy that was obtained for Dataset A. The model makes accurate predictions for 83.4% of the data, up from a chance baseline of 78.3%.

		predicted	
		adverb-final	adverb-initial
observed	adverb-final	4180	5379
	adverb-initial	1955	32659

Table 7: Prediction accuracy for Dataset B, by model fitted on Dataset A

If Dataset B is analyzed on its own terms, we obtain results that are highly similar to those reported for the first model. Again, NP frequency, string resemblance, and collostructional strength emerge as significant predictors. The coefficients are of a similar magnitude as those that were observed in the first model. Table 8 shows the fixed effects of the new model; Table 9 shows its classification accuracy, which is at 83.6% up from a chance baseline of 78.3%.

	Estimate	Std. Error	z value	Pr(>|z|)
Intercept	1.15	0.08	14.86	< 2e-16 ***
NP frequency (logged)	0.36	0.02	17.01	< 2e-16 ***
string resemblance	-1.73	0.15	-11.92	< 2e-16 ***
collostructional strength	0.17	0.04	4.26	2.02e-05 ***

Table 8: Fixed effects in the analysis of Dataset B

		predicted	
		adverb-final	adverb-initial
observed	adverb-final	4135	5424
	adverb-initial	1825	32789

Table 9: Classification accuracy in the analysis of Dataset B

5 Discussion

Our analysis supports the idea of constructional contamination as put forward by Pijpops and Van de Velde (2016): Constructions that are superficially similar influence each other, such that frequent collocations in one construction may influence patterns of variation in another construction. Specifically, we find that higher frequency of an adverb-participle combination in a complex noun phrase increases the likelihood of speakers producing variants of the English passive that are isomorphic to the order of adverb and participle that is found in the noun phrase. As predicted by Pijpops and Van de Velde, constructional contamination is thus a phenomenon that can be observed across different languages. Beyond this general point, our results give rise to three further observations that will be discussed in turn below.

A first point is that our results differ from those of Pijpops and Van de Velde with regard to the variable of string resemblance. Pijpops and Van de Velde (2016: 568) report that high values of string resemblance increase the likelihood that speakers choose a constructional variant that is isomorphic to the contaminating construction; our results indicate the opposite. What accounts for this discrepancy? Possible answers can be gleaned from Figure 1, which shows the distribution of the data with regard to the variables that are at issue, that is, noun phrase frequency, frequency in the adverb-initial passive, and string resemblance.

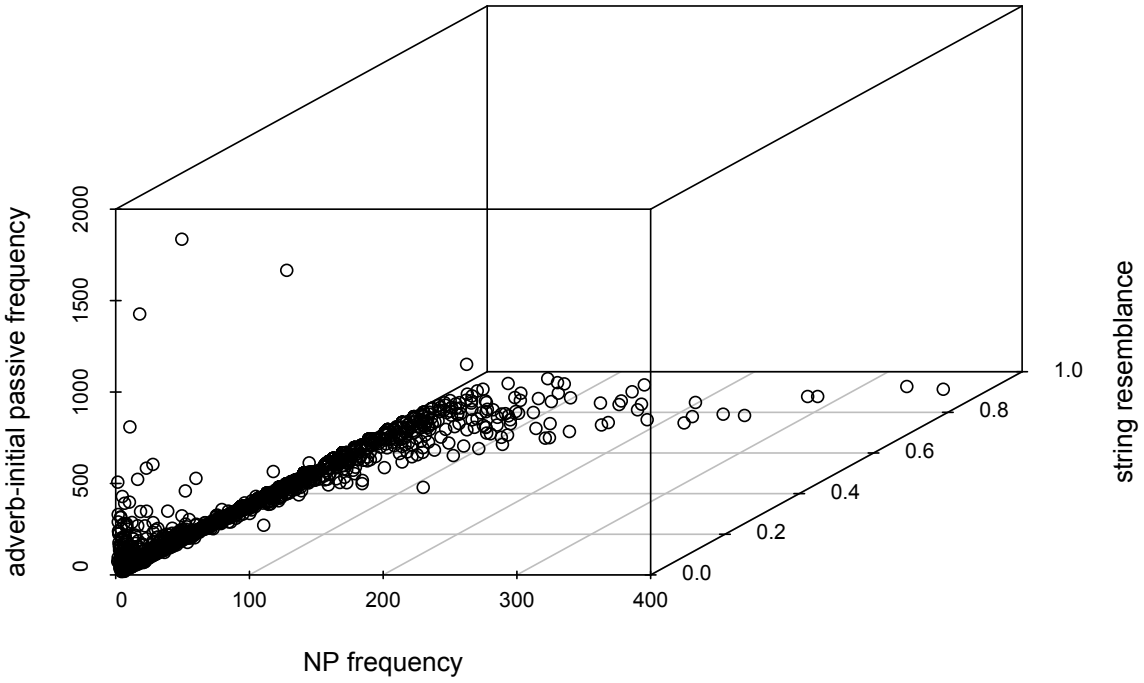

Figure 1: Noun phrase frequency, adverb-initial passive frequency, and string resemblance

The graph shows that most adverb-participle combinations are low in frequency across both the noun phrase construction and the passive construction: There graph shows a heavily overplotted line of data points that are concentrated along the axis where both NP frequency and passive frequency are close to zero. It can also be seen that as string resemblance increases, so does NP frequency: Towards the far end of the cube, there are relatively more points that score relatively high on NP frequency. By contrast, most combinations that are highly frequent in the passive construction are relatively infrequent in

the NP construction and thus score low on string resemblance. We submit that this state of affairs should be rather typical of constructions with high type frequencies. In general, highly frequent constructions, there will always be the tendency that some collocations develop and become strongly conventionalized in only one constructional environment. The constructions that were investigated by Pijpops and Van de Velde were relatively more specific, so that their observations are based on a dataset of 2700 tokens in total (2016: 24). One possible explanation for the divergent results is therefore the size of the dataset, which differs across the two studies. We do not wish to argue that bigger is necessarily better, since the bigger size of our dataset incurs the drawback of diminishing precision and recall due to our reliance on part-of-speech tags. Especially participial forms are quite frequently mistagged. Pijpops and Van de Velde further note that in their multifactorial analysis, the effect of string resemblance is relatively minor (2016: 569). They note that collocations with high string resemblance values are associated with uses of the contaminating construction that are semantically not compatible with the contaminated construction. This is in line with our own observations. Combinations with high string resemblance values include combinations such as *appropriately sized*, *similarly priced*, and *far flung*, which represent fairly unprototypical uses of the English passive. Summing up this point, more studies of constructional contamination with differently-sized datasets will be necessary to come to terms with the effects of string resemblance.

A second point that merits discussion is the fact that constructional contamination can not only be captured through measures that are based on text frequencies of constructions and their variants, but also through collocation measures. Adverb-participle combinations need not be highly frequent, as long as their components are highly informative of each other. In our models, high mutual association of adverb and participle in the noun phrase construction is a significant predictor of speakers' bias towards adverb-initial order in the passive. This observation can be related to findings by Gries, Hampe and Schönefeld (2005: 635), who find that measures of collostructional strength actually outperform text frequency as a predictor of speakers' preferences in production. Our data indicate that text frequency in the noun phrase construction and mutual association of the adverb and the participle are both significant predictors of speakers' choices that explain different aspects of the observed

variance. Taking either predictor out of the regression models leads to a statistically significant impoverishment of the model.

Finally, a third point ties our observations to the general topic of this volume, namely contrast and analogy. Constructional contamination effects can be viewed in terms of contrast and analogical reasoning. Contrast lies at the basis of understanding that a construction such as the English passive may be used in different structural variants. Analogy comes into play when one of these variants is influenced by another construction, such as the complex NP construction. A speaker who has encountered numerous collocations such as *widely believed* in the passive and in the complex NP construction, and who then encounters a new collocation in one of the two constructions, will be able to set up an analogical equation, in which the structural parallelism (Langacker 1999: 145) between the two construction is laid out as follows:

(7) a widely believed rumor : a copiously footnoted paper
---------- = ----------
The rumor was widely believed. : ?

On the basis of previous experience with language use, speakers are able to fill in the missing part of the formula and are subsequently left with a slightly increased bias towards sentences such as *The paper was copiously footnoted*. There is, however, a complication that has been evoked by Boyd and Goldberg (2011) in a study of what they call statistical preemption. Boyd and Goldberg (2011) investigated English adjectives such as *afraid*, *afloat*, or *alive*, which cannot be used attributively (**the afraid child*, **the alive fox*), and which raise the question how speakers acquire this particular constraint. In a nutshell, their answer to this question is that hearers observe a marked asymmetry in speakers' behavior. When speakers use an a-adjective, they systematically avoid a structurally simple attributive construction (**the afraid child*) and instead opt for a more complex relative clause construction such as *the child that was afraid*. Given the frequencies of the attributive adjective construction and of the adjective *afraid*, the absence of their co-occurrence is perceived as significant by hearers. Boyd and Goldberg demonstrate experimentally that speakers obey the non-attributive constraint not only with existing a-adjectives, but also, to

a lesser degree, with novel a-adjectives such as *ablim* or *adax*. This reflects the kind of analogical reasoning that was alluded to above: If *afraid* and *alive* cannot be used attributively, analogy suggests that this is also the case for *ablim* and *adax*.

The complication with regard to analogy and constructional contamination is the following one. In a variant of the original experiment, Boyd and Goldberg used the same experimental set-up but included a training phase in which novel a-adjectives were presented in contexts that prohibited attributive usage for structural reasons. An example of such a context would be *the fox that was ablim and proud of himself* (Boyd and Goldberg 2011: 64). Here, the adjective *ablim* appears in a coordinated adjectival phrase which, due to its complexity, cannot appear in attributive contexts. In other words, the postnominal position of *ablim* is explained by the syntax of its context, and hence there is no need to attribute this behavior to any constraint inherent to the element itself. If participants were exposed to training sentences of this kind, they failed to recognize forms such as *ablim* as belonging to the class of a-adjectives (Boyd and Goldberg 2011: 76). How do these observations relate to the phenomenon of constructional contamination? The main point to take away from Boyd and Goldberg's results is that speakers are selective about applying analogical reasoning. Not all utterances of the same form are processed as potentially analogous to a given construction. More specifically, if a hearer encounters a sequence of adverb and participle, this usage event will only be contaminating the passive if it is actually registered as an instance of the complex noun phrase construction that can be related to the passive. To give a concrete example, the English passive exhibits variance between adverb-initial and adverb-final word order with the combination of *only* and *identified*.

(8) The witness was only identified two weeks later.
The witness was identified only two weeks later.

The string *only identified* does appear in nominal contexts, as in *The only identified witness died in a mysterious car crash*. Would the experience of such an expression trigger a contamination effect? Since *only* in this context functions as an adjective, and since the nominal construction does not relate semantically to the passive in the same way that is outlined in the correspondence in (7), Boyd and Goldberg's results would suggest that such

usage events do not have a contaminating effect. In other words, the mere entrenchment of the string *only identified* is not enough to trigger constructional contamination. The string has to appear in a context that at least potentially alternates with the construction that is to be contaminated. This point is fully in line with the results of Pijpops and Van de Velde, who argue that "both formal and semantic resemblance between constructions is needed to trigger constructional contamination" (2016: 573). What is less clear is what this point implies for the link between constructional contamination and "good enough" parsing (Ferreira et al. 2002, Ferreira and Patson 2007). Pijpops and Van de Velde (to appear) view constructional contamination effects as corroborating evidence for "good enough" parsing, but the semantically restricted nature of these effects could also inspire caution: If speakers operate with chunks of language that are only superficially parsed into syntactic structures, then the mere entrenchment of a sequence of elements actually should yield contamination effects across different syntactic constructions. Apparently, this is not always the case. An important conclusion for constructionist approaches to language is that meaning-based connections between syntactic constructions are at work even during the "shallow" parsing that characterizes early and unproblematic language processing. Discussions of inheritance in Construction Grammar have tended to focus on formal aspects; it seems that the semantic connections in speakers' networks of linguistic knowledge deserve more attention in the future.

6 Concluding remarks

This paper has investigated variation in adverb placement in the English passive, which up to this point has been viewed as a case of largely unpredictable variation. Using the idea of constructional contamination (Pijpops and Van de Velde 2016) as a starting point, we examined whether usage frequencies of English modified noun phrases can explain at least part of the observable variation. Data from the COCA allows us to conclude that this is actually the case. We observe effects of constructional contamination that are based on text frequency and on mutual association between elements that appear in the contaminating construction and the contaminated construction. We linked these observations to a discussion that explored how constructional contamination and analogical reasoning relate

to each other, and we concluded, in line with findings from Boyd and Goldberg (2011) and Pijpops and Van de Velde 2016) that semantic relations between constructions are a crucial precondition for contamination effects to occur.

With an eye to future investigations of constructional contamination, we submit that it still remains to be worked out how corpus frequencies can be operationalized in the best way to measure contamination effects. Our results indicate that simple text frequencies of the contaminating construction are a viable solution, but it goes without saying that there is room for more sophisticated measures such as Pijpops and Van de Velde's string resemblance, which simultaneously takes the frequencies of other constructions into account. It would also be interesting to see further studies that investigate the relative importance of measures based on text frequency and measures based on association strength. Lastly, since semantic relations between constructions have emerged as a decisive factor underlying constructional contamination, it would be rewarding to test whether contamination effects can vary in strength depending on the relative semantic closeness of two or more constructions. Summing up, there is no shortage of follow-up questions, and we look forward to future discussions of these issues.

References

Baayen, R. Harald. (2008). *Analyzing Linguistic Data. A practical introduction to statistics*. Cambridge: Cambridge University Press.

Bates, Douglas, Martin Maechler, Ben Bolker, and Steve Walker (2015). Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software* 67(1), 1-48.

Boyd, Jeremy K., and Adele E. Goldberg. (2011). Learning what not to say: The role of statistical preemption and categorization in *a*-adjective production. *Language* 87, 55-83.

Davies, Mark. (2008). The Corpus of Contemporary American English (COCA): 400+ million words, 1990-present. Available online at http://www.americancorpus.org.

Ferreira, Fernanda, Karl Bailey and Vittoria Ferraro. (2002). Good-enough representations in language comprehension. *Current Directions in Psychological Science* 11(1), 11–15.

Ferreira, Fernanda and Nikole Patson. (2007). The "good enough" approach to language comprehension. *Language and Linguistics Compass* 1, 71–83.

Goldberg, Adele E. (1995), *Constructions. A Construction Grammar Approach to Argument Structure*, Chicago: University of Chicago Press.

Goldberg, Adele E. (2006), *Constructions at Work: The Nature of Generalization in Language*, Oxford: Oxford University Press.

Gries, Stefan Th. (2015). The most underused statistical method in corpus linguistics: Multi-level (and mixed-effects) models. *Corpora* 10(1), 95-125.

Gries, Stefan Th., Beate Hampe, & Doris Schönefeld (2005). Converging evidence: Bringing together experimental and corpus data on the association of verbs and constructions, *Cognitive Linguistics* 16-4, 635-676.

Gries, Stefan Th. and Anatol Stefanowitsch. (2004). Extending collostructional analysis: A corpus-based perspective on "alternations." *International Journal of Corpus Linguistics* 9, 97–129.

Hilpert, Martin. (2014), *Construction Grammar and its Application to English*. Edinburgh: Edinburgh University Press.

Hilpert, Martin. (2015). From hand-carved to computer-based: Noun-participle compounding and the upward-strengthening hypothesis. *Cognitive Linguistics* 26/1, 1-36.

Hilpert, Martin and Holger Diessel. (2017). Entrenchment in construction grammar. In Hans-Jörg Schmid (ed.), *Entrenchment and the psychology of language learning: How we reorganize and adapt linguistic knowledge*. Berlin: Mouton de Gruyter, 57-74.

Huddleston, Rodney and Geoffrey K. Pullum. (2002). *The Cambridge Grammar of the English Language*. Cambridge: Cambridge University Press.

Langacker, Ronald W. (1999). *Grammar and Conceptualization*. Berlin: De Gruyter.

Pijpops, Dirk and Freek Van de Velde. (2016). Constructional contamination: How does it work and how do we measure it? *Folia Linguistica* 50/2, 543-582.

Pijpops, Dirk and Freek Van de Velde. to appear. Constructional contamination in morphology and syntax: four case studies. *Constructions and Frames.*

Stefanowitsch, Anatol & Stefan Th. Gries. (2005). Covarying collexemes. *Corpus Linguistics and Linguistic Theory*, 1(1), 1–43.