Bias against Novelty in Science:
A Cautionary Tale for Users of Bibliometric Indicators

Jian Wang[1], Reinhilde Veugelers[2] and Paula Stephan[3]

[1] University of Leuven

[2] University of Leuven, Bruegel and CEPR

[3] Georgia State University and NBER

December 2015

ABSTRCT

Research which explores unchartered waters has a high potential for major impact but also carries a high uncertainty of having minimal impact. Such explorative research is often described as taking a novel approach. This study examines the complex relationship between pursuing a *novel* approach and impact. We measure novelty by examining the extent to which a published paper makes first time ever combinations of referenced journals, taking into account the difficulty of making such combinations. We apply this newly developed measure of novelty to a set of one million research articles across all scientific disciplines. We find that highly novel papers, defined to be those that make more (distinct) new combinations, have more than a triple probability of being a top 1% highly cited paper when using a sufficiently long citation time window to assess impact. Moreover, follow-on papers that cite highly novel research are themselves more likely to be highly cited. However, novel research is also risky as it has a higher variance in the citation performance. These findings are consistent with the "high risk/high gain" characteristic of novel research. We also find that novel papers are typically published in journals with a lower than expected Impact Factor and are less cited when using a short time window. Our findings suggest that science policy, in particular funding decisions which are over reliant on traditional bibliometric indicators based on short-term direct citation counts and Journal Impact Factors, may be biased against novelty.

Email: jian.wang@kuleuven.be; reinhilde.veugelers@kuleuven.be; pstephan@gsu.edu

Acknowledgements: The authors thank Fabio Montobbio, John Walsh, Jacques Mairesse, Paul David, Henry Sauermann, Jerry Thursby, Stefan Hornbostel, Sybille Hinze, Daniel Sirtes, and Lee Fleming for their very helpful comments. Financial support from KU Leuven (GOA/12/003) and the Research Foundation - Flanders (FWO, G.0825.12) is gratefully acknowledged. J Wang also gratefully acknowledges a postdoctoral fellowship from FWO. Publication data are sourced from Thomson Reuters Web of Science Core Collection.

1. Introduction

Scientific breakthroughs advance the knowledge frontier. Research underpinning breakthroughs often is driven by novel approaches. While novel approaches have a higher potential for major impact, they also face a higher level of uncertainty. In addition, novel research may encounter resistance from incumbent scientific paradigms and consequently suffer from impeded or delayed recognition (Kuhn, 1962; Merton, 1973; Planck, 1950). This “high risk/high gain” nature of novel research makes it particularly appropriate for public support (Arrow, 1962).

There is growing concern, however, that funding agencies are increasingly risk-averse and that their competitive selection procedures encourage relatively safe projects exploiting existing knowledge at the expense of novel projects exploring untested approaches (Alberts, 2010; Kolata, 2009; NPR, 2013; Petsko, 2012; Walsh, 2013). In addition, funding agencies increasingly rely on bibliometric indicators to aid in decision making and performance evaluation (Butler, 2003; Hicks, 2012; Hicks, Wouters, Waltman, de Rijcke, & Rafols, 2015). However, if novel research has an impact profile distinct from non-novel research, using indicators which do not recognize such difference may unintentionally bias funding decisions away from novel research.

In this study, we examine the complex relationship between novelty and impact, using the life-time citation trajectories of 1,056,936 research articles published in 2001 across all scientific disciplines indexed in the Web of Science (WoS), as well as the profile of papers citing them. We demonstrate a “high risk/high gain” profile of novel research, which (i) has a significantly higher probability of leading to major impact directly as well as indirectly (ii) while at the same time has a higher variance in impact performance. We also find (i) that the Impact Factor of the journals in which novel research is published is lower than expected and (ii) that novel papers are less likely to be highly cited in the first few years after publication. Our findings suggest that over-reliance on standard bibliometric metrics, in particular Journal Impact Factors and short-term citation counts, may bias against novel research.

2. Measuring novelty of scientific publications

We view the process of research as one of puzzle solving, whereby researchers work with pieces of knowledge and combine them to generate new scientific knowledge. Using knowledge pieces in well-understood ways corresponds to a search process labeled as *exploitation*. On the other hand, using existing knowledge pieces in new ways corresponds to an *explorative* search process, which is more likely to lead to major breakthroughs but also comes with a substantial risk of no impact (March, 1991). In this perspective, novelty, which is characterized by making new combinations of existing knowledge pieces, is more closely associated with exploration.

This combinatorial view of novelty has been embraced by studies with various disciplinary roots (Burt, 2004; Mednick, 1962; Nelson & Winter, 1982; Schumpeter, 1939; Simonton, 2004), and there is an emerging interest in operationalizing this view to measure the novelty of patents, research proposals, publications, and journals (Boudreau, Guinan, Lakhani, & Riedl, 2014; Fleming, 2001; Lee, Walsh, & Wang, 2015; Packalen & Bhattacharya, 2015; Uzzi, Mukherjee, Stringer, & Jones, 2013; Verhoeven, Bakker, & Veugelers, in press). Following this combinatorial novelty approach, we assess the novelty of a research article by examining the extent to which it makes novel combinations of prior knowledge.

When applying a combinatorial novelty approach to scientific publications, journals can be viewed as representing bodies of knowledge components. Uzzi et al. (2013) have applied this approach to identify papers which make unusual/unexpected/atypical combinations of referenced journals. Rather than looking at the *atypicality* of referenced journal combinations as in (Uzzi et al., 2013), we focus specifically on *novelty* of referenced journal combinations, examining whether an article makes new referenced journal combinations which have never been made in prior publications. Furthermore, we take into account the knowledge distance between the newly-combined journals based on their co-cited journal profiles, i.e., their "common friends." More precisely, we measure the novelty of a paper as the number of new journal pairs in its references weighted by the cosine similarity between the newly-paired journals.

By way of example, a recent breakthrough in human connectome, an automated method to collect serial images of brain tissue for reconstructing the neural network, was accomplished by combining an electron microscope with an ultra-microtome (Denk & Horstmann, 2004). While both knowledge pieces already existed, their combination is a true landmark that revitalized the dream of mapping the human brain (Cook, 2015; Eisenstein, 2009; Kasthuri & Lichtman, 2010). The paper which introduces this novel concept scores in the top 1% of our novelty indicator. Detailed documentation on the construction of the measure is reported in *Appendix I*.

Applying this newly-minted indicator of novelty on a sample of all 2001 WoS research articles, we find that novel research is relatively rare: Only 11% of all papers make at least one new combination of referenced journals. Furthermore, the degree of novelty is highly skewed: More than half (55%) of novel papers make only one new referenced journal pair, typically combining close-by journals, while only a few papers make more new combinations which are typically also more distant (*Table 1*). To further work with this skewedness, we classify papers into three categories: (i) *highly novel*, if a paper has a novelty score among the top 1% of its subject category, (ii) *moderately novel*, if a paper makes at least one new combination but has a novelty score lower than the top 1% of its subject category, and (iii) *non-novel*, if a paper has no new journal combinations.

3. Data

The dataset consists of all research articles published in 2001 indexed in Thomson Reuters Web of Science (WoS). These papers span all scientific disciplines, i.e., 251 WoS subject categories. Our analysis focuses on original research, so our sample consists of only "articles." Other document types, such as "review" and "letter," are excluded. We measure the combinatorial novelty for each article, based on the profile of their references, specifically the newness and unusualness of referenced journal pairs. Therefore, we keep only articles referencing to at least two WoS journals. These are in total 773,311 journal articles, 674,546 of which have at least two references to WoS journals are kept for analysis. We also account for field differences in our analysis; 203 articles with no subject category information are excluded. In addition, 274,072 articles have more than one subject category (up to six subject categories) and are counted multiple times. The final 2001 dataset used has 1,056,936 observations across all 251 WoS subject categories. Citations for each article are also retrieved from WoS in each year from 2001 to 2013. We also look into the profiles of other articles citing our sampled articles, specifically, whether these citing articles are among the top 1% highly cited in the same WoS subject category and publication year.

4. Results

4.1. "High risk/high gain" nature of novel research: Why we (should) care about novelty?

In view of the "high risk/high gain" nature of novel research, we expect novel papers to receive more citations on average but also to have a higher variance in their citation performance. Following (Fleming, 2001), generalized negative binomial models are used to estimate the effects of novelty on both the mean and dispersion of received citations, controlling for potential confounding factors, i.e., field differences, international collaboration, and the number of references and authors (*Fig. 1A* and *Table 3*). We use a 13-year time window to count citations, which is deemed sufficiently long across fields (Wang, 2013). Compared with non-novel papers, *ceteris paribus*, moderately novel papers receive 2% more citations, while highly novel papers receive 11% more citations. At the same time, however, the higher average impact of novel papers is accompanied by a higher variance in citation performance. This higher variance holds most noticeably for highly novel papers, reflecting their higher uncertainty.

A higher variance in impact performance can be driven by more extreme successes and/or more cases of uncited or rarely cited papers. Our analysis shows that highly novel papers are more likely to be among the least 10% cited papers in its field (*Table 3*).

A more important question is whether novel papers are more likely to become "big hits," i.e., receive an exceptionally large number of citations, defined here, following the bibliometric convention, as being top 1% highly cited in the same WoS subject category and publication year. When we use a long (13-year) time window for counting citations, the chance of big hits is more than three times the expected value for highly novel papers and more than 1.5 times for moderately novel papers. Econometric analysis controlling for potential confounding factors confirms this correlation: The odds of big hits is 36% higher for highly novel papers, compared with comparable non-novel papers (*Fig. 1B* and *Table 3*).

Furthermore, novel papers are more likely not only to become big hits themselves but also stimulate follow-up research that leads to breakthroughs. We find that papers that cite novel papers are more likely to themselves receive more citations, compared with papers citing non-novel papers. Likewise, for highly novel papers, the probability of being cited by an article which itself becomes a big hit is about two times the probability for non-novel papers. Econometric analysis, which estimates the probability of a paper being cited by big hits, teasing out any contamination from direct citations received, in addition to controlling for previously mentioned other confounding factors, confirms the positive relation between novelty and indirect impact, (*Fig. 1C* and *Table 3*).

4.2. Bias against novelty: How does novel research score on standard bibliometric indicators?

Given the "high risk/high gain" profile of novel research, the question which follows is: How such research would perform on standard bibliometric indicators. The most influential indicator is probably the Journal Impact Factor. Therefore, we investigate whether novel papers, with their "high risk/high gain" nature, are published in high Impact Factor journals. Although on average novel papers are published in journals with higher Impact Factors, compared with non-novel ones, the multivariate analysis controlling for other confounding factors such as field differences reveals that the Journal Impact Factor of moderately- and highly-novel papers is significantly and substantially lower (approximately 11% and 18% respectively) than comparable non-novel papers (*Fig. 2* and *Table 4*). This finding that novel papers are published in journals with Impact Factors below their potential indicates the obstacles these papers face to being accepted by journals holding central positons in science. Moreover, the negative association between novelty and Journal Impact Factor is not because novel papers are more likely to be published in new journals. Regression analysis additionally controlling for the journal age or whether the journal is new delivers the same conclusion that novel papers are published in lower Impact Factor journals than expected (*Table 4*).

Another widely used bibliometric procedure that we scrutinize is the use of short time windows to count received citations. Despite the fact that novel papers have more citations in the long run as shown in the preceding section, they may take a longer time to be recognized, and such delayed process of citation accumulation might lead to a biased evaluation of impact, when an insufficient time window is adopted to count citations. We estimate the probabilities of being a top 1% highly cited paper for non-, moderately-, and highly-novel papers, using various lengths of time windows (*Fig. 3A* and *Table 5*). In the first few years after publication, novel papers are less likely to be top cited. For example, using a 3-year time window, which is the standard time window used in classic bibliometrics, highly novel papers are not more likely to be a top cited paper than non-novel papers, despite the fact that they have a much higher chance of being top cited when a 13-year citation time window is adopted. As a footnote, although novel papers display a delayed recognition in their direct impact (i.e., being big hits), we find that they have a significantly higher indirect impact (i.e., being cited by big hits) across all citation time windows compared with non-novel papers.

We investigate the bias against novelty and Journal Impact Factors further by additionally incorporating an interaction effect between novelty and Journal Impact Factor, specifically whether the journal in which the focal paper is published has a top 10% Impact Factor in its field. Although publication in a high Impact Factor journal contributes to a faster citation accumulation process, novel papers published in high Impact Factor journals still suffer from a delayed citation accumulation process, compared with comparable non-novel papers in high Impact Factor journals (*Fig. 3B* and *Table 6*). When using a 13-year citation time window, a novel paper in a high Impact Factor journal is more likely to be a big hit than a non-novel paper in a high Impact Factor journal, which is also true when comparing novel and non-novel papers in low Impact Factor journals. However, when using a three-year citation time window, for papers published in low Impact Factor journals, the odds of big hits is 47% higher for highly novel papers than for non-novel papers, but for papers published in high Impact Factor journals, the odds of big hits is 16% lower for highly-novel papers than for non-novel papers. These findings suggest that novel research faces an additional obstacle in the citation accumulation process, even if they successfully enter prestigious journals. This additional obstacle might be explained by the fact that novel research is ahead of its time and needs supporting science and technology to be developed in order to have its potential recognized.

4.3. Robustness analysis

The observed relationship between novelty and impact is universal across fields: Separate analyses by fields yielded similar results (*Appendix II*). We also determined that our results were not sensitive to the procedure of excluding 50% of the least cited journals and requiring that the new combination of journals be reused at least once in the next three years. Relaxing these constraints yielded consistent results. The results are also robust to other variations of our novelty measure, such as excluding referenced star and multidisciplinary journals (*Appendix III*).

Additional analysis also shows that our measure for *novelty* identifies a much rarer phenomenon than the Uzzi et al. (2013) measure for *atypicality* of referenced combinations: While 11% of all publications in our sample score on our *novelty* indicator, 48% make *atypical* combinations. Although both measures of *atypicality* and *novelty* are significantly correlated, they nevertheless are measuring distinct profiles (*Appendix IV*). Papers that score in the top 1% of their field and publication year on *atypicality* are six times more likely to be at the same time highly *novel* papers. Nevertheless, for two thirds of them the *atypical* combinations they make are not *novel*. Furthermore, the impact profile of *atypical* papers is different from *novel* papers. First, the higher dispersion in impact performance plays out more prominently for *novel* papers than for *atypical* papers. Second, *novelty* also has a much stronger effect on stimulating follow-up creativity (i.e., being cited by big hits). Hence, it is not the *atypicality* of combinations being made, but the *novelty* of combinations which identifies the "high risk/high gain" profile of novel research (*Appendix IV*).

5. Discussion

This study contributes to better understanding the relationship between novelty and impact and also uncovers how novel research is "judged" by standard bibliometric indicators. Applying a newly-minted measure of combinatorial novelty on all WoS research articles published in 2001 across all scientific disciplines, we find that novel research (i) has a larger variance in its citation distribution, (ii) is more likely to eventually become a big hit, and (iii) is more likely to stimulate follow-up research which itself become a big hit. These characteristics demonstrate a "high risk/high gain" profile of novel research. However, novel research underperforms in classic bibliometric indicators, (i) being published in journals with lower Impact Factors and (ii) having fewer citations when a short time window is used for counting citations. Therefore, the use of these classic bibliometric metrics would lead to a biased assessment of the value of novel research.

Our findings suggest that caution is called for in using standard bibliometric indicators for funding decisions. Funding agencies are alleged to be increasingly risk-averse and rely on bibliometric indicators for selecting applicants and evaluating their performance. Widely used indicators, such as the Journal Impact Factor and other citation-based metrics using short time windows, run the risk of being biased against novel research given that novel research underperforms on these indicators. This bias imperils scientific progress, as novel research is much more likely to become a big hit itself in the long run as well as to stimulate follow-up big hits. This caveat applies not only to funding decisions but also science policy more generally. The prevailing use of standard bibliometric indicators in various decisions (e.g., hiring and tenure of researchers) at various levels (i.e., department, university, and national) is likely to further disincentivize novel research. We advocate the awareness of such potential bias and suggest, when relying on bibliometric indicators, to use a wider portfolio of indicators. We also call for future research on more nuanced measures to capture other aspects of scientific novelty and explorative research.

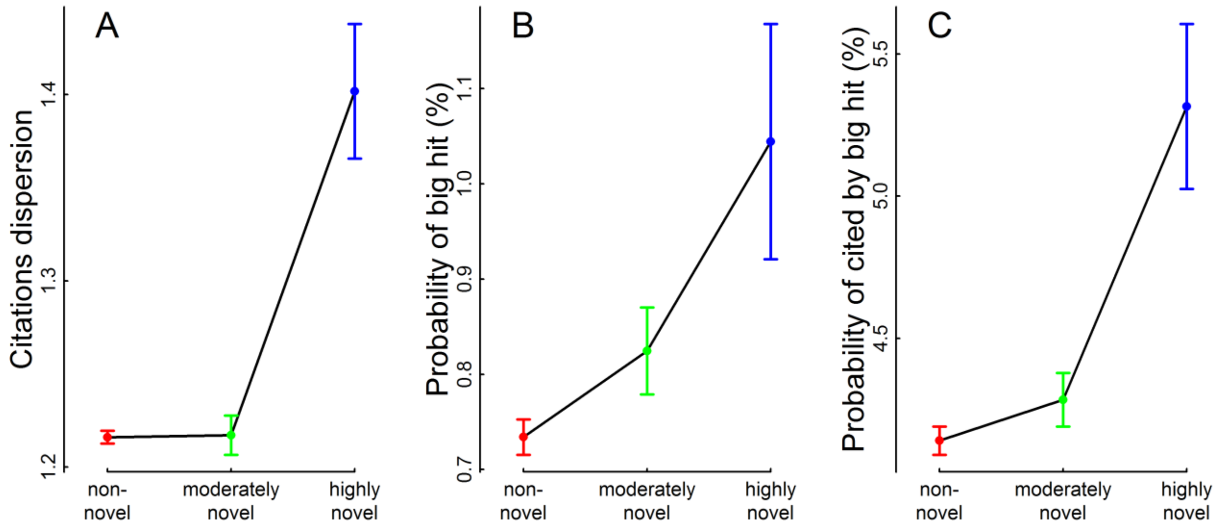

Fig. 1. Impact profile of novel research. (**A**) Estimated dispersion of citations (13-year), using a generalized negative binomial model and controlling for field differences, international collaboration, and the number of authors and references. Estimated values are for an average paper (fixing all other covariates at their means) with different novel classes. The vertical bars represent the 95% confidence interval. Dispersion = (variance – mean) / mean^2. Note that novel papers have both a higher mean and dispersion and therefore also have a much higher variance (variance = mean + dispersion * mean^2). (**B**) Estimated probability of being a big hit, defined as the top 1% most cited articles in each WoS subject category. Results are based on a logistic model. (**C**) Estimated probability of cited by big hits. The citing big hits are identified as the top 1% highly cited articles with the same WoS subject category and publication year, based on their cumulative citations up to 2013. Citing big hits published between 2001 and 2009 are analyzed, and later years are not analyzed because their available time windows are too short to identify big hits reliably. Results are based on a logistic model. All regressions underlying this figure are reported in *Table 3*. Data consist of 1,056,936 WoS articles (no reviews or other document types) published in 2001 and are sourced from Thomson Reuters Web of Science Core Collection.

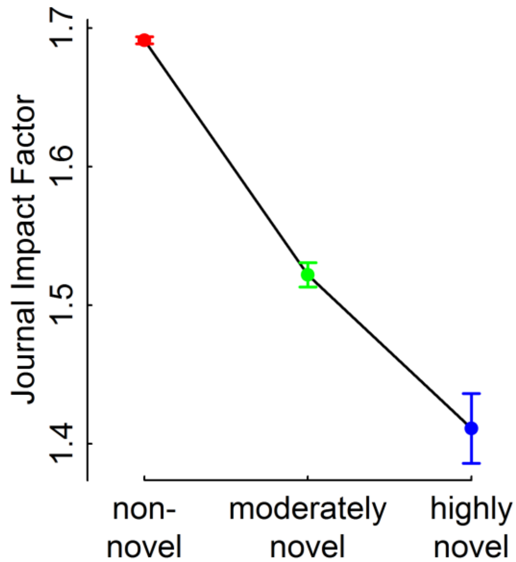

Fig. 2. Journal Impact Factor and novelty. This figure reports the estimated Journal Impact Factor for an average paper with different novelty classes, using a Poisson model and controlling for field differences, international collaboration, and the number of authors and references. Regression outputs in *Table 4*. Data consist of 1,056,936 WoS articles (no reviews or other document types) published in 2001 and are sourced from Thomson Reuters Web of Science Core Collection.

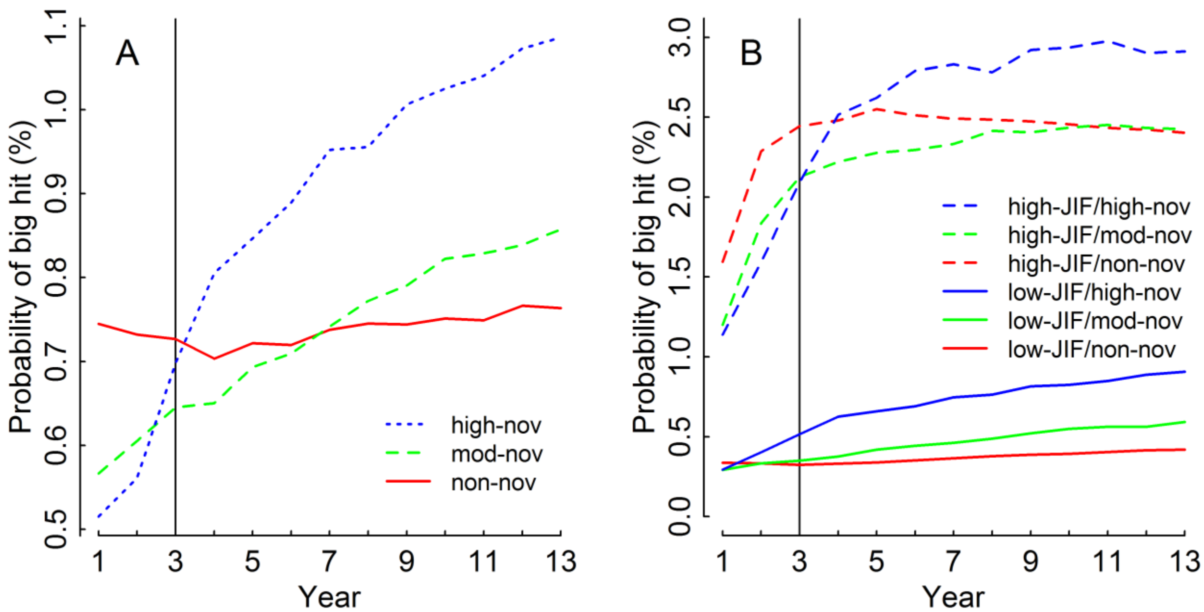

Fig. 3. Citation dynamics and novelty. (**A**) Estimated probability of being a big hit, using 13 consecutive time windows to dynamically identify big hits. As an example, big hits in year 3 are identified as the top 1% highly cited papers based on their cumulative citations in a 3-year time window, i.e., from 2001 to 2003. Results are based on 13 logistic models reported in *Table 5.* (**B**) Estimated probability of being a big hit by year, for papers in different novelty classes and Journal Impact Factor groups. Estimations are based on a set of logistic models additionally incorporating interaction effects between novelty classes and whether a journal has an Impact Factor among the top 10% in its field. Solid lines are estimated probability for papers in low 90% journals with different novelty classes and with all other covariates fixed at their means, and broken lines are for papers in the high 10% journals. Regression outputs are reported in *Table 6.* Data consist of 1,056,936 WoS articles (no reviews or other document types) published in 2001 and are sourced from Thomson Reuters Web of Science Core Collection.

Table 1. The occurrence of novelty

	(1) Number of papers	(2) % of all papers	(3) Avg (avg cos)	(4) Avg(min cos)	(5) avg # new pairs	(6) median # new pairs
non-novel	942850	89%	/	/	/	/
moderately novel	103418	10%	0.0015	0.0013	1.7055	1
highly novel	10668	1%	0.0009	0.0004	7.9499	7

Data sourced from Thomson Reuters Web of Science Core Collection.

We expect our measure to identify only a small minority of papers as novel. Indeed we find that only relatively few papers make new referenced journal combinations. To be more specific, 89% of all papers in our sample do not make any new combinations of referenced journal and therefore do not score on the novelty measure. Of the 11% that make new journal combinations, most (55%) make only one new combination, and only 7% have more than 5 new combinations.

Overall, our measure of novelty displays a highly skewed phenomenon of novelty in scientific publications. To further work with this skewedness in the analysis, we construct a categorical novelty variable *NOV CAT*: (i) *non-novel*, if a paper has no new journal combinations, (ii) *moderately novel*, if a paper makes at least one new combination but has a novelty score lower than the top 1% of its subject category, and (iii) *highly novel*, if a paper has a novelty score among the top 1% of its subject category.

We are particularly interested in papers which are highly (top 1%) novel. These papers make not only more but also more distant new combinations. The median number of new referenced journal pairs they make is 7 (column 6). The new pairs they make are on average more distant, as suggested by the average smaller cosine of their average pairs (column 3), and their most distant pairs have a much higher distance than those made by moderately novel paper, as reflected in the average smaller cosine of the pair with the lowest cosine they make (column 4).

Table 2. List of variables

Variable	Description
Novelty	Combinatorial novelty score.
NOV CAT1	Novelty class dummy: 1 if non-novel, and 0 otherwise.
NOV CAT2	Novelty class dummy: 1 if moderately novel, and 0 otherwise.
NOV CAT3	Novelty class dummy: 1 if highly novel, and 0 otherwise.
JIF	Impact Factor of the journal where the focal paper is published in.
JIF TOP10%	Dummy: 1 if the journal has an Impact Factor among the top 10% in its field.
International	Dummy: 1 if internationally co-authored, and 0 otherwise.
Authors	The number of authors.
Refs	The number of references.
C_t	Cumulative number of citations in year *t*, i.e., 2001-*t*, where $t \in [2001, 2013]$.
Big hit in year *t*	Dummy: 1 if among the top 1% highly cited in year *t*, based on C_t, in the same WoS subject category and publication year. $t \in [2001, 2013]$.
Cited by big hits	Dummy: 1 if cited by a big hit published between 2001 and 2009. The citing big hits are identified as the top 1% highly cited article in the same WoS subject category and publication year, based on their cumulative citations in 2013. Only big hits published between 2001 and 2009 are identified, so that each citing article has at least five years to accumulate citations, for a reliable identification of citing big hit articles.

Table 3. "High risk/high gain" nature of novel research

	Citations (13-year)		Least 10% cited (13-year)	Top 1% cited (13-year)	Cited by big hits ('01-'09)
	generalized negative binomial		logit	logit	logit
	mean	dispersion			
	(1)	(2)	(3)	(4)	(5)
NOV CAT2	0.0241***	0.0009	-0.0664***	0.1171***	0.0356**
	(0.0036)	(0.0047)	(0.0163)	(0.0286)	(0.0110)
NOV CAT3	0.1116***	0.1421***	0.1102*	0.3557***	0.2620***
	(0.0108)	(0.0132)	(0.0559)	(0.0602)	(0.0292)
9-year citations (ln)					1.8363***
					(0.0051)
International	0.0755***	-0.0680***	-0.2357***	0.0593*	0.0037
	(0.0028)	(0.0036)	(0.0127)	(0.0240)	(0.0089)
Authors (ln)	0.2755***	-0.1437***	-0.6315***	0.5754***	-0.0648***
	(0.0020)	(0.0024)	(0.0076)	(0.0181)	(0.0066)
Refs (ln)	0.6331***	-0.2559***	-1.1328***	1.2209***	-0.1143***
	(0.0020)	(0.0024)	(0.0067)	(0.0207)	(0.0071)
N	1020561		866544	1056895	1056565
Log lik	-4114822		-218464	-56275	-266709
Chi2	223598***		42288***	5543***	145503***

Data consist of all WoS articles published in 2001 and are sourced from Thomson Reuters Web of Science Core Collection.
Field (WoS subject category) fixed effects incorporated.
Clustered-robust standard errors in parentheses.
*** p<.001, ** p<.01, * p<.05, + p<.10.

For the generalized negative binomial model (column 1-2), WoS subject categories with fewer than 1000 papers are excluded for reliable estimates of the dispersion parameter. While the original sample has 1,056,936 articles of 251 subject categories, regressions reported here use 1,020,561 articles of 170 subject categories.

Cited by big hit (column 5) means cited by an article, which is published in between 2001 and 2009 and among the top 1% highly cited articles within the same WoS subject category and publication year, based on their cumulative citations up to 2013. Big hits published after 2009 are not analysed because their available time windows are too short to identify big hits reliably.

Table 4. Journal Impact Factor and novelty

	JIF Poisson	JIF Poisson	JIF Poisson
NOV CAT2	-0.1055***	-0.1033***	-0.0839***
	(0.0031)	(0.0031)	(0.0031)
NOV CAT3	-0.1811***	-0.1744***	-0.1352***
	(0.0092)	(0.0091)	(0.0088)
Journal age < 4		-0.3421***	
		(0.0048)	
Journal age (ln)			0.2211***
			(0.0013)
International	0.0737***	0.0724***	0.0651***
	(0.0024)	(0.0023)	(0.0023)
Authors (ln)	0.1735***	0.1712***	0.1602***
	(0.0018)	(0.0018)	(0.0017)
Refs (ln)	0.3519***	0.3496***	0.3250***
	(0.0017)	(0.0017)	(0.0016)
N	1056936	1056936	1056936
Log pseudo-likelihood	-1716396	-1711208	-1680856
Chi2	942910***	966429***	1169732***

Data consist of all WoS articles published in 2001 and are sourced from Thomson Reuters Web of Science Core Collection.
Field (WoS subject category) fixed effects incorporated.
Clustered-robust standard errors in parentheses.
*** p<.001, ** p<.01, * p<.05, + p<.10.

Results of Poisson models are reported here, an alternative specification, using the OLS model and the log of JIF as the dependent variable yields consistent results.

Table 5. Big hits and novelty

	Big hit ’01-‘01 logit	Big hit ’01-‘02 logit	Big hit ’01-‘03 logit	Big hit ’01-‘04 logit	Big hit ’01-‘05 logit	Big hit ’01-‘06 logit	Big hit ’01-‘07 logit	Big hit ’01-‘08 logit	Big hit ’01-‘09 logit	Big hit ’01-‘10 logit	Big hit ’01-‘11 logit	Big hit ’01-‘12 logit	Big hit ’01-‘13 logit
NOV	-0.2754***	-0.1919***	-0.1196***	-0.0790*	-0.0412	-0.0147	0.0051	0.0358	0.0613*	0.0916**	0.1021***	0.0908**	0.1171***
CAT2	(0.0398)	(0.0332)	(0.0314)	(0.0306)	(0.0299)	(0.0296)	(0.0294)	(0.0292)	(0.0289)	(0.0287)	(0.0287)	(0.0288)	(0.0286)
NOV	-0.3709***	-0.2682**	-0.0405	0.1366*	0.1607*	0.2133**	0.2572***	0.2510***	0.3048***	0.3136***	0.3315***	0.3397***	0.3557***
CAT3	(0.1043)	(0.0828)	(0.0715)	(0.0655)	(0.0641)	(0.0626)	(0.0619)	(0.0623)	(0.0609)	(0.0610)	(0.0606)	(0.0604)	(0.0602)
Internat	0.2324***	0.1973***	0.1400***	0.1181***	0.1002***	0.1036***	0.0992***	0.0913***	0.0927***	0.0946***	0.0877***	0.0674**	0.0593*
ional	(0.0280)	(0.0245)	(0.0240)	(0.0239)	(0.0237)	(0.0237)	(0.0237)	(0.0238)	(0.0238)	(0.0238)	(0.0238)	(0.0240)	(0.0240)
Authors	0.6041***	0.7648***	0.8207***	0.7973***	0.7659***	0.7485***	0.7249***	0.6932***	0.6613***	0.6363***	0.6139***	0.5877***	0.5754***
(ln)	(0.0209)	(0.0192)	(0.0187)	(0.0185)	(0.0181)	(0.0181)	(0.0180)	(0.0179)	(0.0179)	(0.0180)	(0.0180)	(0.0181)	(0.0181)
Refs	0.9759***	1.2277***	1.2835***	1.2977***	1.3013***	1.2958***	1.2779***	1.2583***	1.2491***	1.2380***	1.2282***	1.2259***	1.2209***
(ln)	(0.0241)	(0.0220)	(0.0215)	(0.0213)	(0.0208)	(0.0208)	(0.0208)	(0.0207)	(0.0206)	(0.0207)	(0.0206)	(0.0207)	(0.0207)
N	1052801	1055915	1056516	1056757	1056778	1056574	1056703	1056803	1056895	1056781	1056813	1056902	1056895
Log lik	-41211	-50133	-52495	-53555	-54675	-55026	-55249	-55582	-55858	-55849	-56074	-56209	-56275
Chi2	3753***	5871***	6745***	6815***	6919***	6858***	6604***	6300***	6195***	5967***	5780***	5613***	5543***

Data consist of all WoS articles published in 2001 and are sourced from Thomson Reuters Web of Science Core Collection. Field (WoS subject category) fixed effects incorporated. Clustered-robust standard errors in parentheses. *** p<.001, ** p<.01, * p<.05, + p<.10.

Big hits in year t are identified as the top 1% highly cited papers within the same WoS subject category, based on their citations received between 2001 and t. Identifying big hits using annual citation counts instead of cumulative citation counts yields similar results.

Table 6. Big hits, novelty interacts with JIF

	Big hit ’01-‘01 logit	Big hit ’01-‘02 logit	Big hit ’01-‘03 logit	Big hit ’01-‘04 logit	Big hit ’01-‘05 logit	Big hit ’01-‘06 logit	Big hit ’01-‘07 logit	Big hit ’01-‘08 logit	Big hit ’01-‘09 logit	Big hit ’01-‘10 logit	Big hit ’01-‘11 logit	Big hit ’01-‘12 logit	Big hit ’01-‘13 logit
NOV CAT2	-0.1436**	0.0021	0.0765	0.1255**	0.2129***	0.2310***	0.2382***	0.2540***	0.2993***	0.3354***	0.3315***	0.3068***	0.3502***
	(0.0550)	(0.0495)	(0.0479)	(0.0463)	(0.0443)	(0.0435)	(0.0429)	(0.0423)	(0.0412)	(0.0406)	(0.0403)	(0.0403)	(0.0396)
NOV CAT3	-0.1364	0.1931+	0.4661***	0.6375***	0.6671***	0.6769***	0.7210***	0.7068***	0.7489***	0.7463***	0.7449***	0.7654***	0.7781***
	(0.1438)	(0.1138)	(0.0989)	(0.0899)	(0.0877)	(0.0859)	(0.0837)	(0.0836)	(0.0814)	(0.0814)	(0.0807)	(0.0793)	(0.0787)
JIF TOP10%	1.5661***	1.9481***	2.0416***	2.0346***	2.0388***	1.9851***	1.9432***	1.9035***	1.8743***	1.8537***	1.8140***	1.7846***	1.7683***
	(0.0274)	(0.0252)	(0.0249)	(0.0247)	(0.0245)	(0.0243)	(0.0242)	(0.0241)	(0.0240)	(0.0240)	(0.0239)	(0.0238)	(0.0239)
JIF TOP * NOV CAT2	-0.1441+	-0.2263**	-0.2189***	-0.2386***	-0.3294***	-0.3228***	-0.3052***	-0.2834***	-0.3278***	-0.3437***	-0.3232***	-0.3016***	-0.3411***
	(0.0780)	(0.0658)	(0.0626)	(0.0609)	(0.0592)	(0.0584)	(0.0579)	(0.0573)	(0.0568)	(0.0563)	(0.0561)	(0.0563)	(0.0558)
JIF TOP * NOV CAT3	-0.2052	-0.5628**	-0.6241***	-0.6226***	-0.6392***	-0.5684***	-0.5894***	-0.5909***	-0.5779***	-0.5627***	-0.5374***	-0.5791***	-0.5816***
	(0.2072)	(0.1631)	(0.1409)	(0.1288)	(0.1262)	(0.1233)	(0.1219)	(0.1228)	(0.1203)	(0.1204)	(0.1197)	(0.1196)	(0.1193)
International	0.1901***	0.1488***	0.0907***	0.0707**	0.0535*	0.0586*	0.0554*	0.0489*	0.0515*	0.0543*	0.0484*	0.0289	0.0216
	(0.0280)	(0.0246)	(0.0241)	(0.0240)	(0.0238)	(0.0238)	(0.0239)	(0.0239)	(0.0239)	(0.0239)	(0.0240)	(0.0241)	(0.0241)
Authors (ln)	0.4802***	0.6222***	0.6751***	0.6507***	0.6193***	0.6033***	0.5813***	0.5501***	0.5190***	0.4942***	0.4737***	0.4481***	0.4366***
	(0.0216)	(0.0199)	(0.0195)	(0.0193)	(0.0190)	(0.0189)	(0.0188)	(0.0188)	(0.0188)	(0.0188)	(0.0189)	(0.0189)	(0.0189)
Refs (ln)	0.7393***	0.9478***	0.9939***	1.0154***	1.0211***	1.0235***	1.0117***	0.9960***	0.9924***	0.9846***	0.9803***	0.9821***	0.9806***
	(0.0257)	(0.0242)	(0.0239)	(0.0237)	(0.0232)	(0.0231)	(0.0230)	(0.0229)	(0.0228)	(0.0228)	(0.0227)	(0.0227)	(0.0227)
N	1052801	1055915	1056516	1056757	1056778	1056574	1056703	1056803	1056895	1056781	1056813	1056902	1056895
Log lik	-39508	-46753	-48568	-49580	-50638	-51165	-51537	-51998	-52395	-52481	-52831	-53062	-53209
Chi2	8514***	14380***	16198***	16406***	16649***	16089***	15574***	15093***	14746***	14305***	13933***	13624***	13310***

Data consist of all WoS articles published in 2001 and are sourced from Thomson Reuters Web of Science Core Collection. Field (WoS subject category) fixed effects incorporated. Clustered-robust standard errors in parentheses. *** p<.001, ** p<.01, * p<.05, + p<.10.

Journals are classified into two groups: the top 10% in the same WoS subject category and others. Using a different classification, the top 25% versus the rest yields consistent results.

References

Alberts, B. (2010). Overbuilding Research Capacity. *Science, 329*(5997), 1257-1257. doi: 10.1126/science.1197077

Arrow, K. (1962). Economic welfare and the allocation of resources for invention *The rate and direction of inventive activity: Economic and social factors* (pp. 609-626): Princeton University Press.

Boudreau, K. J., Guinan, E., Lakhani, K., & Riedl, C. (2014). Looking Across and Looking Beyond the Knowledge Frontier: Intellectual Distance and Resource Allocation in Science. *Available at SSRN 2478627*.

Boyack, K. W., & Klavans, R. (2014). Atypical combinations are confounded by disciplinary effects. *STI 2014 Leiden*, 64.

Burt, R. S. (2004). Structural holes and good ideas. *American Journal of Sociology, 110*(2), 349-399. doi: 10.1086/421787

Butler, L. (2003). Explaining Australia's increased share of ISI publications—the effects of a funding formula based on publication counts. *Research Policy, 32*(1), 143-155. doi: 10.1016/s0048-7333(02)00007-0

Cook, G. (2015, January 11). Sebastian Seung's Quest to Map the Human Brain. *The New York Times Magazine,* MM27.

Denk, W., & Horstmann, H. (2004). Serial block-face scanning electron microscopy to reconstruct three-dimensional tissue nanostructure. *Plos Biology, 2*(11), 1900-1909. doi: 10.1371/journal.pbio.0020329

Eisenstein, M. (2009). Putting neurons on the map. *Nature, 461*(7267), 1149-1154. doi: 10.1038/4611149a

Fleming, L. (2001). Recombinant uncertainty in technological search. *Management Science, 47*(1), 117-132. doi: 10.1287/mnsc.47.1.117.10671

Hicks, D. (2012). Performance-based university research funding systems. *Research Policy, 41*(2), 251-261. doi: 10.1016/j.respol.2011.09.007

Hicks, D., Wouters, P., Waltman, L., de Rijcke, S., & Rafols, I. (2015). The Leiden Manifesto for research metrics. *Nature, 520*, 429-431.

Kasthuri, N., & Lichtman, J. W. (2010). Neurocartography. *Neuropsychopharmacology, 35*(1), 342-343. doi: 10.1038/npp.2009.138

Kolata, G. (2009, June 27, 2009). Grant System Leads Cancer Researchers to Play It Safe, *New York Times*. Retrieved from http://www.nytimes.com/2009/06/28/health/research/28cancer.html

Kuhn, T. S. (1962). *The structure of scientific revolutions*. Chicago, IL: University of Chicago Press.

Lee, Y.-N., Walsh, J. P., & Wang, J. (2015). Creativity in scientific teams: Unpacking novelty and impact. *Research Policy, 44*(3), 684-697. doi: http://dx.doi.org/10.1016/j.respol.2014.10.007

March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization Science, 2*(1), 71-87. doi: 10.1287/orsc.2.1.71

Mednick, S. A. (1962). The associative basis of the creative process. *Psychological Review, 69*(3), 220-232. doi: 10.1037/h0048850

Merton, R. K. (1973). *The sociology of science*. Chicago, IL: University of Chicago Press.

Nelson, R. R., & Winter, S. G. (1982). *An Evolutionary Theory of Economic Change*. Cambridge, MA: Belknap Press of Harvard University Press.

NPR. (2013). Scientists win Nobel for work on how cells communicate.   Retrieved April 29, 2015, from http://www.npr.org/templates/story/story.php?storyId=230192033

Packalen, M., & Bhattacharya, J. (2015). Neophilia Ranking of Scientific Journals: National Bureau of Economic Research.

Petsko, G. A. (2012). Goodbye, Columbus. *Genome Biology, 13*(5). doi: 10.1186/gb-2012-13-5-155

Planck, M. (1950). *Scientific autobiography, and other papers*. London,: Williams & Norgate.

Schumpeter, J. A. (1939). *Business Cycles; A Theoretical, Historical, and Statistical Analysis of the Capitalist Process* (1st ed.). New York, London,: McGraw-Hill Book Company, inc.

Simonton, D. K. (2004). *Creativity in Science: Chance, Logic, Genius, and Zeitgeist*. Cambridge, UK; New York, NY: Cambridge University Press.

Uzzi, B., Mukherjee, S., Stringer, M., & Jones, B. (2013). Atypical combinations and scientific impact. *Science, 342*(6157), 468-472. doi: 10.1126/science.1240474

Verhoeven, D., Bakker, J., & Veugelers, R. (in press). Measuring Technological Novelty with patent based indicators. *Research Policy*.

Walsh, D. (2013). Not Safe for Funding: The N.S.F. and the Economics of Science. *The New Yorker*. Retrieved April 29, 2015, from http://www.newyorker.com/tech/elements/not-safe-for-funding-the-n-s-f-and-the-economics-of-science

Wang, J. (2013). Citation time window choice for research impact evaluation. *Scientometrics, 94*(3), 851-872. doi: 10.1007/s11192-012-0775-9

Appendix I. Novelty measure

We develop the following procedure to operationalize the combinatorial novelty measure.

- For each focal paper, we retrieve all of its co-cited journal pairs (J_1-J_2, J_1-J_3, J_1-J_4 …)
- We check each journal pair to see if it is new, i.e., never appeared since 1980. The 1980 cut off is because of data-availability reasons. It assumes a window of 20 years before obsolescence.
- For those new journal pairs (e.g., J_1-J_2), we check how easy it is to make this new combination by looking at how many common "friends" they have. More specifically, we compare the co-citation profiles of the two journals (J_1 and J_2) in the preceding three years (between year t-3 and t-1).
    - We use the following matrix where each row or column provides the co-citation profile for a journal. The i,j-th element in this symmetric matrix is the number of times that J_i and J_j are co-cited, that is, the number of papers published between year t-3 and t-1 that cite them together. For example, in the preceding three years, J_1 and J_2 have never been cited together by any papers, but J_1 and J_3 have been cited together by 3 papers.

	J_1	J_2	J_3	J_4	J_5	…
J_1	/	0	3	0	5	…
J_2	0	/	6	2	3	…
J_3	3	6	/	5	4	…
J_4	0	2	5	/	0	…
J_5	5	3	4	0	/	…
…	…	…	…	…	…	/

    - The ease of combining J_1 and J_2 is defined as the cosine similarity between their co-citation profiles:

$$COS_{1,2} = \frac{J_1 \cdot J_2}{\|J_1\|\|J_2\|}$$

where J_1 and J_2 are row (or column) vectors. Cosine similarity is a classic measure of similarity between two vectors, and is widely used in bibliometrics.

- Correspondingly, the difficulty of combining J_1 and J_2 is: $1 - \mathrm{COS}_{1,2}$.

- For each paper, we construct a continuous indicator of combinatorial novelty as the sum of all new combinations weighted by the cosine-based ease of making the new combination. Papers without new combinations get 0 by definition.

$$Novelty = \sum\nolimits_{J_i - J_j \text{ } pair \text{ } is \text{ } new} \left(1 - COS_{i,j}\right)$$

- To avoid trivial combinations, we focus only on the most important journal combinations, i.e., we exclude 50 percent of the least cited journals (as measured in the preceding three years). To further reduce the likelihood of picking up trivial combinations, we impose as a condition that the new combination has to be reused at least once in the next three years. We check the robustness of the main results to these choices in *Appendix II & III*.

Appendix II. Robustness: Scientific fields

The inclusion of scientific field dummies in the econometric analysis corrects for field specific effects influencing impact but does not allow for any field specificity in the relationship between novelty and impact. For example, are novel papers more likely to lead to big gains in some fields, while other fields are more averse to novel papers, hampering their impact?

To examine this in more detail we perform an analysis of the main effects of novelty by scientific discipline. We use 3 groups: LS (Life Sciences), PSE (Physical Sciences and Engineering), and SS (Social Sciences). In LS, we distinguish LS2 (Medicine) from the rest (LS1). In PSE we distinguish PSE2 (Computer Sciences; Engineering) from the rest. Both LS2 and PSE2 are the more applied counterparts of LS1 and PSE1. PSE2 and especially SS are relatively small fields compared with LS1, LS2 and PSE1, which may hamper significant effects for these fields.

The results, available on request from the authors, show that the finding that novel papers are less likely to be published in high impact journals holds especially in LS1, LS2, and PSE1. But in PSE2 and SS, the negative association between novel paper and the Journal Impact Factor is insignificant. The result on the higher dispersion of citations for novel papers holds for all subfields, with the exception of the SS, where there is no significant difference in the dispersion of impact for novel and non-novel papers. In all subfields, highly novel papers are more likely to be top 1% cited papers when using the long citation time window. When using a shorter (3-year) citation time window, and not controlling for the Journal Impact Factor, in none of the scientific fields are novel papers significantly more likely to be top 1% highly cited. Only when controlling for the Journal Impact Factor is there a significant positive effect for highly novel papers, but only in LS1 and PSE1.

We conclude that the main results are more or less robust by field, especially for the larger fields LS and PSE. The result on the higher variance in impact for novel papers and the higher likelihood to be top 1% cited with a longer time window are robust across all fields. The effect of the lower Journal Impact Factor is less robust in the smaller fields.

Appendix III. Robustness: Indicator variations

The *novelty* indicator used in the analysis excluded the 50% least cited journals and required reuse of the new combination in the next three years. To check the sensitivity of the analysis to these choices, we replicate the analysis without these restrictions. Although most new journals are already not included in the analysis because they are typically among the lower cited journals, we also checked the results excluding new and/or young journals for the calculation of *novelty*. In addition, Boyack and Klavans (2014) warned that the *atypicality* measure in (Uzzi et al., 2013) is confounded with citing star journals, such as *Science*, *Nature*, and *PNAS*, although they did not directly test whether such confounding effect would change the findings of (Uzzi et al., 2013). Our *novelty* measure is unlikely to be confounded with citing star journals, which are so highly cited that combining them with any other journals is less likely to be new. Nevertheless, we check results excluding top 10% highly cited journals and multidisciplinary journals.

We test robustness of our findings using the following variations of the *novelty* measure:

- V1. Only exclude 50% least cited journals
- V2. Only require reuse in 3 years
- V3. Exclude new journals
- V4. Exclude journals younger than 3 years
- V5. Exclude 50% least cited journals, require reuse in 3 years, and exclude top 10% highly cited journals
- V6. Exclude 50% least cited journals, require reuse in 3 years, exclude top 10% highly cited journals, and exclude multidisciplinary journals

The result on the significantly lower journal impact factor for novel papers is found across all specifications. The result on novelty being associated with a significantly higher dispersion of citations is robust and similar in size across all specifications. The higher probability of being in the top 1% highly cited papers, when taking a long citation time window (13 years) is also robust across specifications.

If we include all cited journals and only impose the reuse requirement, we find a higher probability for highly novel papers to be in the top 1% highly cited papers when using a 3-year window, even without correcting for the JIF. This is likely because the inherent association between reuse and citations is more visible when not excluding least cited journals. In contrast, if we do not impose the reuse requirement but only exclude the least cited or youngest journals, we get a negative association between novelty and being top 1% cited when taking the short 3-year citation window and ignoring the JIF effect. Even correcting for the JIF effect we find no significant positive associations.

Appendix IV. Relationship with the Uzzi measure for atypical combinations

We compare our novelty measure with the Uzzi (Uzzi et al., 2013) measure. Following (Uzzi et al., 2013), we construct an *atypicality* measure (using their terminology) for a paper based on its left tail of the distributions of referenced-journal-pairs typicality values. Specifically, we follow the adapted version in (Lee et al., 2015). The commonness for each referenced journal pair is first calculated and sorted, and then the 10th percentile of this distribution of commonness value at the journal-pair level gives the *atypicality* score at the paper level. In a similar manner to our categorical *NOV CAT* measure, we define a categorical *atypicality* measure *UZZ ATYP*, which takes a score of 3 for papers which are among the top 1% of their subject category on Uzzi's *atypicality* measure, a score of 1 for papers without any atypical journal combinations (i.e., the observed number of co-citations is smaller than the expected). We assign a score of 2 for everything else. *Table A1* reports the two-way table between *NOV CAT* and *UZZ ATYP*.

Table A1: Co-occurrence between *novelty* and *atypicality*

	NOV CAT=1 (89%)	NOV CAT=2 (10%)	NOV CAT=3 (1%)
UZZ ATYP=1 (52%)	1.06	0.49	0.34
	94.9%	4.7%	0.4%
	54.8%	25.1%	17.6%
UZZ ATYP=2 (47%)	0.94	1.52	1.61
	83.5%	14.9%	1.6%
	44.4%	72.2%	76.6%
UZZ ATYP=3 (1%)	0.75	2.75	5.75
	67.3%	26.9%	5.8%
	0.8%	2.8%	5.8%

For each contingency, we report in black the ratio of observed to expected frequency (expected in the case of independency between *novelty* and *atypicality*), in blue the row percentages (i.e., the distribution across *NOV CAT* by *UZZ ATYP*), and in red the column percentage (i.e., the distribution across *UZZ ATYP* by *NOV CAT*). Data sourced from Thomson Reuters Web of Science Core Collection.

One first observation from *Table A1* is that a lower share of papers are defined as *novel*, using our measure, than are defined as *atypical*, using the Uzzi scoring method: While only 11% of the papers in our sample score on *novelty*, 48% of papers score on the Uzzi *atypicality* measure, i.e., almost half of all papers are *atypical*. Our measure therefore picks up a rarer novelty phenomenon, than does *atypicality*.

Table A1 confirms a significant positive association between *atypicality* and *novelty*. The χ^2 test of independence between row and column variables is highly significant ($\chi^2 = 2491$, $p < 2.2e\text{-}16$). Being *non-novel* but *atypical* (cells 4 & 7) and being *non-atypical* but *novel* (cells 2 & 3) occur less frequently than expected, while being *non-novel* and *non-atypical* (cell 1) and being *novel* and *atypical* (cells 5, 6, 8 and 9) occur more frequently than expected. Especially noteworthy is that scoring in the top 1% range on both *atypicality* and *novelty* occurs 5.75 times more often than expected. Despite this positive association, *Table A1* also shows that the off-diagonal cases of non-overlap are substantial in numbers:

- 45.2% of all papers with no novel combinations (i.e., *NOV CAT* = 1) score positively on *atypicality* (cells 4&7 in column 1). Their *atypical* combinations are, however, not *novel*.
- Of the top 1% atypical papers (i.e., *UZZ ATYP = 3*), two thirds do not make new combinations (cells 8 & 9 in row 3);
- Although 82.4% of the *highly novel* papers (i.e., *NOV CAT* = 3) also make on average atypical combinations (cells 6&9 in column 3), only 5.8% of them score in the top 1% on *atypicality*, meaning that although they are making substantially new combinations, the profile of all the combinations they make is only moderately atypical.

The comparison confirms that although *novelty* and *atypicality* are related, they are nevertheless distinct concepts. *Novelty* may be one of the drivers for *atypicality* of combinations, but *atypicality* is not a direct measure of *novelty*, capturing a less skewed phenomenon.

We further run regressions, adding the Uzzi measure and compare its effects against our novelty measure. *UZZ ATYP* and *NOV CAT* have different distributions across their three categories, which might influence the comparison results. Therefore, in addition to *UZZ ATYP*, we construct an alternative categorization for Uzzi's *atypicality*, with the same proportion of papers in each category as our novelty categories. Specifically, *UZZ CAT* = 3 if top 1% in the same WoS subject category and publication year, 2 if below top 1% but above top 10%, and 1 all others. In addition, we also compare Uzzi's *atypicality* and our *novelty* score, both continuous variables. We find that:

- Both *novelty* and *atypicality* have positive effects on the dispersion of 13-year citations, but *novelty* has a more pronounced effect.
- Both *novelty* and *atypicality* have positive effects on the probability of big hits when using a 13-year citation time window. *Atypicality* seems to have a larger effect than *novelty*. However, when we correct the disadvantage of novel papers being published in lower Impact Factor journals (additionally control for the Journal Impact Factor the regressions), *novelty* has a larger effect than *atypicality*.
- Both *novelty* and *atypicality* have positive effects on the probability of cited by big hits, but *novelty* has a more pronounced effect.
- *UZZ CAT* has a negative effect on JIF, but this negative effect disappears after adding *NOV CAT* into the regression. On the other hand, *UZZ ATYP* and *atypicality (ln)* do not have negative effects on JIF.
- While *novelty* has a negative effect on the probability of big hits using a three-year time window, *atypicality* does not.
- In all cases, the effects of *novelty* remains consistent, with or without having *atypicality* in the regression at the same time.

We conclude that our novelty measure captures a rarer phenomenon than Uzzi's atypicality measure. While both measures are significantly correlated, they are also sufficiently distinct. The impact analysis of our novelty measure is robust when we additionally control for atypicality. Furthermore, the impact profile for *NOV CAT* shows more pronounced effects with respect to higher risk and delayed impact compared with the atypicality, dimensions which are more closely associated with novel research.