

**Killer Incentives: Rivalry, Performance and
Risk-Taking among German Fighter Pilots, 1939-45**

<i>Philipp Ager</i>	<i>Leonardo Bursztyn*</i>	<i>Lukas Leucht</i>	<i>Hans-Joachim Voth</i>
University of Southern Denmark & CEPR	University of Chicago & NBER	UC Berkeley Haas School of Business	University of Zurich & CEPR

This Version: June 2021

Abstract. Using newly-collected data on death rates and aerial victories of more than 5,000 German fighter pilots during World War II, we examine the effects of public recognition on performance and risk-taking. When a particular pilot is honored publicly, both the victory rate and the death rate of his *former* peers increase. Fellow pilots react more if they come from the same region of Germany, or if they worked closely with him. Our results suggest that personal rivalry can be a prime motivating force, and that non-financial rewards can lead to a crowd-in of both effort and risk-taking via social connections.

Keywords: Non-financial rewards, rivalry, status competition, World War II.

* Corresponding author (bursztyn@uchicago.edu). For helpful comments we thank the Editor (Adam Szeidl), five anonymous referees, Chris Blattman, Karol Borowiecki, Thomas Dohmen, Ruben Durante, Armin Falk, Ernst Fehr, Raquel Fernandez, Bruno Frey, Mitch Hoffman, Jonathan Leonard, Adriana Lleras-Muney, Ricardo Perez-Truglia, Luigi Pascali, Michael Peters, Konrad Raff, Gautam Rao, Debraj Ray, David Stasavage, Noam Yuchtman, as well as seminar audiences at the AEA meetings, Academy of Behavioral Economics, Bonn, Harvard, NBER Summer Institute, NYU, Northwestern, Stanford, UC Berkeley, UC Berkeley Haas, UCLA, Southern Denmark, and Zurich. We are indebted to Veit Scherzer for help with German World War II data, to Iftach Spector and Zeev Raz of the IDF for insightful discussions, and to Jim Perry, Tony Wood, and Johannes Mathews for sharing their data. We thank Vanessa Sticher, Victoria Gierok, and Jacob Miller for outstanding research assistance.

I. INTRODUCTION

Humans are social animals; it is part of human nature to compare ourselves with others. From Adam Smith (1759) to Veblen (1899) and Duesenberry (1949), economists have argued that individuals care about their *relative* position. Many institutions and firms exploit such concerns: prizes are handed to the best teacher, artist, or athlete; firms reward workers for relative performance, and professors assign grades based on class rank. Empirical evidence suggests that recipients of non-pecuniary rewards show lower absenteeism, greater effort, and higher accuracy.¹ At the same time, a “medal glitters, but it also casts a shadow”, causing “disappointments and heartburnings” (Winston Churchill):² every act of recognition also lowers the relative standing of those *not* recognized. Lower rank often has negative consequences: higher earnings by neighbors correlate with lower self-reported happiness; knowledge of relative salary can affect job satisfaction, subjective well-being, and choice of the city of residence.³

Despite recent advances in the literature on relative standing concerns, several important questions remain open. First, how do these concerns operate in higher-stakes settings? Indeed, much of the evidence comes from relatively low-stakes settings. Job satisfaction in a survey and overspending on platinum credit cards are not life-or-death outcomes; as with many experimental results, the external validity of (field) experiments with low stakes can be questioned. Second, mechanisms are often unclear. Relative standing can serve an instrumental function: to be “top” can lead to dating success or the best house in town. Alternatively, they can reflect private, intrinsic concerns about relative standing. Beyond that, who are people comparing themselves with? Are relative standing concerns driven by one's perceptions of one's overall position, or can these concerns stem from personal comparisons, including personal rivalries? To answer these questions, one needs a real-world setting which also allows for the separation of these two mechanisms. In addition, ruling out time-variant confounders requires high-frequency outcomes.

In this paper, we examine a high-stakes setting where high-frequency outcomes and personal rivalry can be observed – aerial combat during World War II. We demonstrate the

¹ Markham et al. (2002); Chan et al. (2014); Kosfeld and Neckermann (2011); Ashraf et al. (2014). Neckermann et al. (2014) find that employees continue to perform strongly after they received an award. The opposite emerges in the work of Malmendier and Tate (2009), who show that superstar CEOs significantly underperform after receiving major awards.

² Knowles (1999, p. 215).

³ See for example Luttmer (2005); Card et al. (2012), Perez-Truglia (2020); Bottan and Perez-Truglia (2020). In a laboratory setting, Kuziemko et al. (2014) show that individuals go to great lengths to avoid being ranked last in a group. These studies and our work relate to the psychology literature on social comparisons, going back to the seminal work by Festinger (1954).

importance of personal competition and rivalry as a key motivating force behind relative standing concerns. Our setting allows us to study behavioral responses to exogenous changes in an individual's relative position within a highly personal reference group – and without clear instrumental benefits. Specifically, we examine the effects of peer recognition on risk-taking and performance, using newly assembled data on the death rates and aerial ‘victory’ scores of German fighter pilots.⁴

Aerial combat is a useful setting for analyzing what motivates highly skilled individuals: The stakes are high, the task complex, and both death and performance are well-measured. Crucially, once an air battle begins, there is no effective control of individual planes by superior officers. In every ‘dogfight,’ the pilot has to decide whether to continue the attempt to shoot down the enemy plane, or to break off contact. Motivating top pilots was key for aggregate performance; a carefully designed system of medals and other awards created incentives for pilots to “keep scoring.” Top pilots mattered for aggregate outcomes: Overall, the 5,000 pilots in our dataset downed 55,000 Allied planes but performance was highly skewed. The best 1% (5%) in our dataset accounted for 15% (43%) of all victories, destroying 8,500 (23,800) Allied aircraft. At the same time, over four fifths (4,100) of pilots in our data died, were captured, or invalidated out.⁵

We focus on one type of public recognition – mentions in the German armed forces daily bulletin (*Wehrmachtbericht*). Armies have long used “mentions in dispatches” to recognize particular achievements. The German daily bulletin typically contained a summary of military developments. Occasionally, it would highlight individual accomplishments, such as a high number of enemy ships sunk, of tanks destroyed, or a “round” number of cumulative air victories, like 150 or 200.⁶ Mentions constituted an exceptional form of recognition and were as rare as the most prestigious medals. The bulletin was distributed widely – it was broadcast on the radio, published in the press and on command posts throughout German-held territory. No simple rule “entitled” a soldier to being mentioned. Moreover, unlike other settings in which individuals compete for an award or recognition (such as in sports competitions), there was no fixed number of mentions that pilots were competing for.⁷

⁴ Our use of the terms ‘victory’, ‘score’ or ‘ace’ are not meant to glorify the deeds of German fighter pilots, who served a genocidal regime. The use of these terms is standard in historical research about military pilot performance in World War I and World War II.

⁵ Actual performance was even more concentrated at the top – to enter our dataset, a pilot had to have filed at least one victory claim. The majority of pilots never did. While luck undoubtedly played a large role, skill differences are necessary to explain overall inequality (Simkin and Roychowdhury 2006).

⁶ We draw on Wegmann (1982), an edited compendium of all *Wehrmachtbericht* issues.

⁷ As a result, if the pilots were interested in (the relatively small) instrumental benefits associated with a public mention, responses to peer mentions were driven by relative concerns and not by changes in the perceived scarcity of future mentions.

Mentions were highly visible to others for a brief moment; in contrast to high-ranked medals, they did not create a permanently visible sign of distinction.

To identify effects, we focus on the recognition of a *former peer* (a pilot who used to fly in the same squadron in the past). During their service, pilots were assigned to squadrons, the primary fighting unit of every air force, consisting of 8-12 pilots. Frequent transfers meant that as time went by, they acquired a growing group of former fellow flyers. When one pilot was recognized in a highly visible way for his accomplishments, his former peers showed large increases in both death and victory rates. In normal times, pilots in our data (who at some point have a peer mentioned in the daily bulletin) ‘score’ 0.7 victories and die at a rate of 2.7% per month. When a former peer is mentioned, however, these rates jump by more than 50% – to 1.4 victories per month, and a death rate of over 4.1%. We confirm these results by creating placebo mentions – pilots who were not recognized but *could* have been because they performed similarly well as actually mentioned pilots. We do so by following an approach similar to Azoulay et al. (2019), using Coarsened Exact Matching (CEM) to create placebo mentions for high-performing pilots. Reassuringly, the former peers of these placebo-mentioned pilots do not react after a fictitious mention.

Public recognition of *any pilot* may impact a pilot’s *overall* standing in the air force, and increase tangible rewards: promotions, leave, and other possible benefits in case of German victory may have depended on a pilot's relative standing in the air force as a whole. Recognition of a *former peer* however does not change a pilot’s future benefits from improving his rank in the air force as a whole: it only diminishes his relative standing in a *personally defined* peer group of (former) comrades. We control for performance changes in response to any pilot being mentioned, and focus on the additional effect of a former peer receiving recognition. This *relative decline* in personal standing due to a former peer being recognized is largely unobserved by others.

Consistent with this interpretation, we show that pilots do not react more strongly to a peer’s recognition when their social “universe” (the set of pilots with whom they flew in the past) contains more former peers of the mentioned pilot – the reaction is personal, and not mediated by the size of the common peer group. We also find that pilots do not react more strongly when another pilot from their same birthplace region gets mentioned – *unless the mentioned pilot was a former peer*. This helps rule out a mechanism based on social image concerns vis-à-vis individuals back in their home regions, or on the competition for instrumental benefits at the local level. Taken together, our results indicate that the changes in risk-taking and performance in response to a former peer’s recognition are likely to reflect an

intrinsic, *personal* urge to restore one's relative standing. This need was sufficiently strong to affect life-and-death decisions. Our findings highlight the importance of spillovers in social groups, and the personal nature of incentives: recognition of an individual can lead to very different responses by others, depending on the nature of their personal connection to that individual.

How do these effects arise? Pilots could either put in more effort, flying more missions – or they could try harder during each flight, taking more risks. Using a newly collected dataset based on individual pilots' personal flight logbooks, we show that pilots fly more combat missions after a peer was recognized. This suggests that the provision of extra effort is an important mechanism. In addition, we examine the drivers behind the change in the performance of pilots. We find that pilots ran a higher risk of death per combat day, and for each victory. This implies that they also increased their risk-taking, going after “marginal” and more dangerous and difficult victories immediately after a peer had been mentioned.

Our paper contributes to a growing literature on image concerns and behavior. Recent papers show that individuals care about how others view them and that this has important effects on an important range of behaviors, from charitable donations (DellaVigna et al. 2012), to campaign contributions (Perez-Truglia and Cruces 2017), voting (DellaVigna et al. 2017), protest participation (Enikolopov et al. 2020), credit card take-up (Bursztyn et al. 2018), and educational investments (Bursztyn and Jensen 2015). Compared to these studies, we evaluate the potential role of image concerns in a setting with extremely high stakes. We also show that an important part of the underlying mechanism is less about social image *vis-à-vis others*, as in the literature on “keeping up with the Joneses” in terms of consumption. Bertrand and Morse (2016) provide evidence that the consumption of the median household in a state is predicted by variation in the income of the top quartile. Kuhn et al. (2011) and Agarwal et al. (2016) show that neighbors of lottery winners are more likely to face financial distress.

Instead, we show that it can have a direct link to measurable output, with *personal rivalry* a driving force. Our paper complements the existing literature on “keeping up with the Joneses” in a number of ways: i) we examine a high-stakes setting beyond the standard settings of consumption and income generation; ii) our analysis sheds light on who people compete with and provides evidence that relative position concerns can be very personal; iii) we also provide evidence of a specific psychological mechanism: in a high-stakes setting, our

findings indicate that these relative concerns matter even in the absence of clear instrumental benefits.⁸

We also add to the literature on tournaments. There are strong theoretical grounds for believing that – in a single-shot setting – tournaments can induce greater effort from participants (Lazear and Rosen 1981; Green and Stokey 1983; Nalebuff and Stiglitz 1983a; 1983b). However, many tournaments are dynamic in nature; the step-by-step release of information in such a setting has the potential to transform incentives in important ways (Lizzeri et al. 2002; Yildirim 2005; Ederer 2010; Goltsman and Mukherjee 2011); for empirical evidence on differential risk-taking in tournaments see, for example, Brown (2011), Fershtman and Gneezy (2011), and Genakos and Pagliero (2012). Our own results indicate that status concerns can indeed promote risk-taking, and we demonstrate this dynamic in a setting with high stakes (and no tangible upside, financially). Far from pilots “giving up”, though, we find that rivalry can lead to a crowding-in of effort and greater risks taken – with deadly consequences.⁹

Studies on peer effects examine how collaborating with others affects worker effort and performance (Falk and Ichino 2006; Bandiera et al. 2010). There is some well-identified evidence in modern settings that workers’ effort increases when exposed to more productive peers (Mas and Moretti 2009), and that the strength of social ties mediates such peer effects (Bandiera et al. 2005). Shue (2013) demonstrates that compensation and acquisitions policies correlate more among executives who were part of the same (randomly formed) section of MBA students, and that class reunions heighten this correlation. Relatedly, overrepresentation of one’s own gender and race in randomly-assigned groups can reduce promotion prospects (Karaca-Mandic et al. 2013). Relative to that literature, we make three contributions. First, we provide insight on mechanisms. Because we focus on information about *former* peers, knowledge spillovers and task complementarities can be ruled out as drivers of our findings, highlighting the importance of social comparisons in our setting. Second, we show that these social spillovers can be important in high-skill occupations. This is important because there is so far only limited evidence of peer effects among the highly skilled (Jackson and Bruegmann 2009; Azoulay et al. 2010; Waldinger 2012). Third, we find high performance under relative incentives. While we do not observe an alternative incentive structure overall (such as piece rates), we do find that the stronger relative status concerns, the higher performance was. This

⁸ The “keeping up with the Joneses” literature generally cannot separate intrinsic from instrumental drivers, such as benefits in the housing and dating markets.

⁹ Relatedly, the literature on status (Besley and Ghatak 2008; Moldovanu et al. 2007; Chan et al. 2014; Frey 2007) has shown how recognition can sharpen workplace incentives.

contrasts sharply with Bandiera et al. (2005), who find markedly lower performance under relative incentives.

Finally, we add to the literature on conflict participation. How to structure incentives in a military setting is important – precisely because the standard tools of personnel economics are likely to be blunt (Acemoglu et al. 2020; Costa and Kahn 2003). While our findings refer to a specific setting and group of individuals, every army – and many non-military organizations – use non-financial incentives to create status competition amongst employees (Frey 2007; Besley and Ghatak 2008; Moldovanu et al. 2007). Military data on pilots allows us to use individual-level outcome data and to measure risk-taking. We find that high-powered incentives – in the form of public recognition – may backfire precisely because concerns about relative standing can induce too much risk-taking.¹⁰ Whereas the literature on incentives in the military has traditionally emphasized the importance of collaborative effort (Stouffer et al. 1949; McPherson 1997; Van Creveld 2007), we highlight the importance of individual incentives and personal competition as key determinants of military performance.

The paper proceeds as follows. Section II provides background on the German air force during World War II and on the data we use. In Section III we present the main findings, and Section IV discusses mechanisms, alternative interpretations and presents additional robustness checks. We conclude in Section V.

II. HISTORICAL BACKGROUND AND DATA

In this section we describe the setting of our study: The organization of the German air force in World War II and its rise and fall as a fighting force. We also discuss the sources and limitations of our data.

A. The German air force during World War II

The German air force entered World War II with around 4,000 planes, including 1,200 fighters, and 880,000 men (Kroener et al. 1988). In the opening phase of the war, the German air force achieved air supremacy. The only exception before 1943 was the defeat during the Battle of Britain. The planned invasion of the British Isles had to be called off because of the Luftwaffe’s failure to dominate the skies. By 1943, both personnel and the number of planes had approximately doubled (Kroener et al. 1988). As the Allied bomber offensive against German cities gathered pace, ever more fighter units were called back to defend the German

¹⁰ One clear analogy is bonuses in financial institutions, where, the desire to be the “best” trader or loan officer can lead to catastrophic economic losses. Brown et al. (1996) argue that relative performance incentives can lead to excessive risk-taking in asset management.

Reich. Having started the war with modern planes and a large air fleet, Germany first lost its quantitative edge. Once the Wehrmacht invaded Russia, and the United States joined the war, the Luftwaffe was heavily outnumbered in all theaters of war. It eventually fell behind also in terms of equipment quality. New planes with advanced technology, such as the ME-262 jet, arrived too late to make a difference. Pilot training also suffered: by 1944, a typical German pilot accumulated less than half the flying hours of UK and U.S. pilots before being sent into combat.

Loss rates increased over the course of the war, reaching staggering levels. By January 1942, the German air force lost 1.8% of its fighter pilots; by May 1944, it was losing 25% of them every month (Evans 2009). The destruction of planes was even more rapid: The Luftwaffe lost 785 planes in combat (and another 300 in accidents, etc.) during the six months between June and November 1940; between January and June 1944, it lost 2,855 aircraft in combat, plus another 1,345 in accidents (Murray 1996). Despite these losses, high production rates ensured that the actual number of fighters in combat units continued to rise until the end of 1944. Nonetheless, by 1944 the Luftwaffe was a much-diminished force. German air force pilots by then had less training than Allied pilots, fuel for aircraft was often in short supply, and the Luftwaffe was numerically outnumbered in the East and in the West.

B. Rank and public recognition

Aerial victories are the key determinant of relative standing among fighter pilots. Attaining ‘ace’ status is an important concern as highlighted in the memoirs of surviving pilots from WWI to Yom Kippur. James Salter, a US fighter pilot during the Korean War, wrote in an autobiographical novel: “**If you did not have [victories], you were nothing.**”¹¹ Heinz Knoke, for example, an ace with 33 confirmed victories, first flew with the 52nd fighter wing (*Jagdgeschwader*). The unit mainly operated in the East, and many high-scoring aces served in it. After being transferred to the West in 1942, Knoke (1953) noted in his diary:

Some of my old comrades in the No. 52 Fighter Wing have been killed in action in recent months. The score of enemy aircraft shot down to their credit has mounted [...] Barkhorn and Rall are credited with having shot down more than 100 Russian planes. It is a keen disappointment for me not to have been able to stay with the old gang.¹² (p. 54)

¹¹ “[The aces] stood out like men moving forward through a forest of stumps. Their names were gilded. They had shot down at least five [enemy planes]. [...] There were no other values [...] That was the final judgement. **[Victories] were everything. ...If you did not have [victories], you were nothing.**” (Salter 1956, emphasis added).

¹² Similarly, during the Battle of Britain – arguably the decisive air battle of World War II – two highly-decorated German pilots, Adolf Galland and Werner Mölders, were neck and neck in terms of total victories (Galland 1993). When Mölders was ordered to confer with the head of the *Luftwaffe*, Hermann Göring, he went

Knoke was not only keeping track of his former teammates' overall score. Despite several deaths of comrades, he reacted in a strong, personal way to the widening gap between himself and his former squadron peers. In our main analysis, we will use peer groups at the same level – former squadron peers – as mentioned by Knoke to define joint service.

Recognition for aerial victories took two main forms – medals and mentions. The German armed forces operated an elaborate system of medals. The principal awards for valor were the Iron Crosses and the Knight's Cross, with higher awards requiring increasingly higher tallies of downed Allied fighters.¹³ In addition, soldiers could receive a mention in the daily bulletin. This was one of the highest forms of recognition available in the German armed forces. A typical daily report would describe battles on the different fronts. Mentions of individual soldiers' 'accomplishments' were rare: During the entire war, fewer than 1,200 men were recognized in this way (Wegmann 1982), out of the 18 million German men who served.¹⁴ A typical example for fighter pilots is Hans-Joachim Marseille's second mention on June 18, 1942: "First Lieutenant Marseille shot down ten enemy planes in a 24-hour period in North Africa, raising his total score of aerial victories to 101" (Wegmann 1982).

The propaganda department within the operations staff of the German armed forces produced the daily *Wehrmachtbericht*. Like all propaganda by the Third Reich, it mixed truth and distortions (Scherzer 2005). Highlighting the alleged "superiority" of German fighting men was an integral part of this strategy. However, we find no evidence of the *Wehrmachtbericht* distorting the accomplishments of pilots. Mentions only occur for an exclusive group of outstanding pilots. We have information on the universe of mentioned fighter pilots (60); of these, 43 are mentioned for the number of aerial victories they achieved, either cumulatively or in a single period (one day, one month, etc). Mentioned pilots ended the war with an average of 107 victories, and scored an average of 2.6 victories a month (compared to an average of 0.67 victories per pilot-month in our overall sample).

C. Data

Our database of German fighter pilots during World War II draws on two principal sources: Jim Perry and Tony Wood's *Oberkommando der Luftwaffe* (OKL) combat claims list, and the

to Berlin for three days of meetings – but only on the condition that Galland be grounded for the same number of days. Göring (himself a high-scoring WWI fighter pilot) agreed, taking one of his best pilots out of combat for no reason other than to ensure 'fair' competition.

¹³ During World War II, about 3.3 million Iron Crosses 2nd class were awarded but only 7,300 Knight's Crosses, 890 Knight's Crosses with Oak Leaves, 160 with Swords, 27 with Diamonds, and one with Golden Oak Leaves.

¹⁴ There are 1,182 individual surnames in the *Wehrmachtbericht*. Because first names are not always recorded, there could be as many as 1,739 soldiers mentioned (if each mention with an identical last name is for a different soldier).

Kracker Luftwaffe Archive.¹⁵ The OKL fighter claims list was extracted from microfilms of the handwritten records of the Luftwaffe personnel office (Personalamt) stored at the German Federal Archives (Bundesarchiv) in Freiburg. Because some OKL fighter claims records did not survive the war, Tony Wood augmented the list with claims from other published sources – such as Donald Caldwell’s (1996) JG 26 war diary – to obtain a comprehensive list of German fighter claims for the years 1939-1945.

We construct a monthly panel based on a cleaned version of the Perry-Wood fighter claims, aggregating the information for every pilot by month and year. This panel contains the number of monthly victories per pilot together with pilots’ first and last name, rank, wing, group, and squadron. We then match the panel data with additional information from the Kracker Luftwaffe Archive. Kracker’s archive contains detailed personal data on German fighter pilots, collected from several sources, such as their war status (e.g., killed in action, prisoner of war, World War II survivor), and for some pilots also the starting date of his Luftwaffe career. Thus, for every pilot in the sample, we have information on their monthly victories, their squadron unit, whether he received an award, and how long he was active during World War II. We combine this with information on combatant status (i.e., MIA, POW, and whether they were killed or wounded). Our database does not include pilots who never shot down a plane during aerial combat.

We only analyze the behavior of daytime fighter pilots. This is because the tasks and skills of day and night fighter pilots differed substantially. Whereas day fighters mainly battled against other fighter pilots, night fighters were mainly used to intercept bombers (Murray 1996). Our sample is unbalanced and consists of 5,081 fighter pilots of the German Luftwaffe that made at least one combat claim during World War II. Pilots are observed for 16 months on average, for a total of 82,348 observations. Next, we augment our information with death dates of pilots taken from pilot biographies (Mathews and Foreman 2015). These biographies are based on primary sources, principally microfilms from the Bundesarchiv in Germany and unit war diaries. We additionally incorporate information on pilots that exit our sample after getting shot down and wounded from lists compiled by Matti Salonen and from the *Luftwaffe Officer Career Summaries* edited by Henry L. deZeng IV and Douglas G. Stankey.¹⁶ In our data, we find that of the 5,081 pilots, 4,109 (or 80.9%) exit the sample –

¹⁵ For more information about Tony Wood’s combat claims list and the Kracker Luftwaffe Archive, see <https://web.archive.org/web/20130928070316/http://lesbutler.co.uk/claims/tonywood.htm> and <http://www.aircrewremembered.com/KrackerDatabase/>.

¹⁶ While Mathews and Foreman (2015) only publish biographies of pilots with at least five claims, we are grateful to Johannes Mathews for sharing with us his 7,730 biographies of pilots with at least one claim. When merging the biography data into our data set based on pilot names, we get 3,932 exact matches. Additionally, we manually went through possible matches proposed by probabilistic matching and confirmed 128 of them as correct. Our final biographical data contains details on 4,060 out of 5,081 pilots. We are grateful to Matti

meaning they are not in the next month's data set (provided the war has not yet ended). 2,531 out of 4,109 exits in our data are documented in at least two of our sources. For 2,724 of the exits we can confirm the death of a pilot. Our sources also refer to some of the other exiting pilots as getting shot down, being taken as prisoners of war, missing in action, or being severely wounded. This suggests that the vast majority of cases indeed refer to pilots who were either killed or permanently incapacitated.¹⁷

The high command of the German air force (*Oberkommando der Luftwaffe*, OKL) received fighter claims throughout the war. A special staff for recognition and discipline was in charge of collecting and validating claimed aerial victories. Pilots were required to file extensive documentation before a claim was recognized. The OKL records contain information on every reported aerial victory of German fighter pilots during World War II by wing (*Geschwader*), unit (*Gruppe*), squadron (*Staffel*), and pilot's name and rank as well as by the day, location (grid reference), type of damage, witnesses, and type of the claimed aircraft. German rules for counting a claim as an aerial victory were demanding; enemy aircraft crash sites on German territory were routinely investigated (Caldwell 2012). Each claim had to be accompanied by a witness report confirming either the destruction of the enemy plane (impact or explosion in the air), or that the enemy pilot had bailed out. Many claims were not accepted, – for example, Heinz Knoke was credited with 33 victories, but claimed an additional 19 Allied planes without receiving confirmation by the OKL.¹⁸

Some pilots accumulated a large number of 'victories.' The highest-scoring fighter pilot in history was Erich Hartmann, with 352 confirmed claims (Mathews and Foreman 2015). Primary determinants of victory rates were experience, luck, skill, and the number and type of opposing Allied planes. Leading German 'aces' recorded higher numbers of victories than Allied pilots largely because German pilots did not have a pre-specified "tour" after which they would rotate out of active service – the rule was "fly till you die" (Spick 2011). Altogether, German air force records document 54,835 confirmed destructions of Allied aircraft by pilots in our dataset. In an average month, the average German pilot 'scored' 0.67 victories and faced a 5% risk of exiting the sample permanently (which was practically synonymous with death). In the East (West), the victory rate was 1.15 (0.39) and the exit rate

Salonen for sharing his complete lists on plane crashes in which pilots were wounded. 3681 of our pilots could be matched to this list. For more information on the *Luftwaffe Officer Career Summaries* see <http://www.ww2.dk/lwoffz.html>, which we accessed on April 1st, 2019. We matched 1927 of our pilots to the career summaries.

¹⁷ To be conservative, we always refer to "exit rates" instead of "death rates" when describing our statistical results in the remainder of this paper.

¹⁸ We do not use unconfirmed victories. While some evidence suggests a tendency to "over-claim" by both the Western and German air forces (Caldwell 2012), support for underclaiming also exists (McFarland and Newton 2006). This has probably less to do with systematic dishonesty and more with the confusion of air combat (Galland 1993).

0.043 (0.054). In other words, the exchange ratio (the number of Allied planes shot down before a pilot was lost) was 27 in the East and 7 in the West.¹⁹

The distribution of the number of Allied planes shot down is highly uneven. In an average month, the vast majority (almost 80%) of pilots failed to score even a single victory. At the same time, some pilots quickly notched up large numbers of victories: Emil Lang shot down 68 Allied planes in October 1943, and Hans-Joachim Marseille scored 17 victories in a single day (September 1, 1942). The top-scoring 350 pilots achieved more aerial victories than the 4,700 lowest-scoring pilots combined. Appendix Figure A.1 graphs the number of monthly victories per pilot by the quantiles of the distribution.

There was a large seasonal component to air combat. The summer season – when ground operations were common and hours of daylight were long – saw spikes in aerial activity; the winter months brought a lull in fighting. Appendix Figure A.2 plots the mean victory and exit rates over time. The time-series peaks mostly coincide, except for the end of the war when the victory rate plummeted and the exit rate increased.

D. Organization and training

The German air force was divided into air fleets (*Luftflotten*), each of which was responsible for a particular geographical area. The number of fleets rose from four to seven during World War II. Air corps within each air fleet controlled the planes and men; air “districts” were responsible for infrastructure. The air corps consisted of wings (*Geschwader*) of 100-150 planes each. The wings were organized by function, with different *Geschwader* for fighter planes, long-range bombers, dive bombers, reconnaissance, and so forth. Each wing typically comprised three groups (*Gruppe*) each consisting, in turn, of three or four squadrons (*Staffel*). Every squadron had an authorized strength of twelve aircraft, but the actual number could be as high as sixteen or as few as four or five aircraft (Stedman and Chappell 2002).

Air force units were not recruited on a regional basis. There is also no evidence that the better graduates from the air combat schools were sent to elite squadrons. The allocation of new pilots to units was largely random, driven by operational needs, recent losses, and – sometimes – personal connections in addition to transfers for disciplinary reasons (Caldwell 1996). Andrew (2004) argues that there was a “lack of formal controls systems” in the Luftwaffe and attributes the German air force’s decline and fall to its lack of formalized procedures and standardized management routines.²⁰ For example, the transfer of Hans-

¹⁹ This is *not* the standard definition of the exchange ratio, which normally measures either planes for planes or pilots for pilots. Here we calculate the number of enemy planes shot down in exchange for every pilot lost.

²⁰ Individual squadron histories give many examples of transfers and new pilots being sent into units to make up losses. Given that up to a quarter of all Luftwaffe pilots died each month (USSBS 1945), the urgent need to replenish squadrons was a key determinant of personnel assignment.

Joachim Marseille from the Channel Front to North Africa was meant to punish him for frequently turning up for duty hung over and overtired (Kurowski 1994). The squadron history of JG 26 recorded 13 arrivals and seven departures in a single month; arrivals contained ex-members transferring back, as well as rookies, while those rotating out were sent to other squadrons, instrument training, or instructor duty (Caldwell 1996).

III. MAIN EMPIRICAL RESULTS

In this section, we demonstrate our main finding: Former peers of a pilot who received public recognition – in the form of a mention in the daily armed forces bulletin – substantially increase performance. They also take greater risks, causing them to exit our sample at a higher rate, which means death in most cases. Pilot skill modifies these patterns – “top” pilots mostly reacted to a peer’s recognition by improving their scores, whereas average pilots were more likely to die. We first document these patterns with a simple survival analysis and fixed effects models in the full sample of all pilots. We build on this preliminary evidence by implementing an event-study framework in a trimmed sample of similar pilots based on Coarsened Exact Matching (following Iacus et al. 2012): Both exit and victory rates increase sharply after a former peer’s mention. In addition, we show that pilots are more likely to react strongly to the mention of a former peer if they were born in close geographical proximity, or if they served together in smaller, more tightly-knit units.

A. Performance and risk-taking of former peers

We are interested in whether pilots exit at a higher rate and score more victories after one of their peers is mentioned. Over time, pilot performance and exit rates are strongly correlated within squadrons since pilots in the same unit are subject to similar shocks (weather, enemy activity, etc.). To sidestep the reflection problem, we focus on the effect of *former* peers being mentioned in the Wehrmacht bulletin.

Figure 1 illustrates our identification strategy, using the case of two pilots – Franz Doerr and Walter Schuck. Doerr was mentioned in June 1944, for scoring 12 ‘victories’ within a single day. Both pilots flew with the 7th squadron of JG 5 until October 1942; Doerr then continued to serve there until July 1944, whereas Schuck was transferred to the 9th squadron of JG 5 in October 1942. At the time of Doerr’s mention, he had not flown with Schuck for 20 months. We now examine the performance of Schuck in June 1944, when Doerr was mentioned. *In this subsection*, we compare his victory and exit rates with those of other pilots who do not have a former peer mentioned during the same month. Some 1% of our observations refer to pilot-months when a former peer of a pilot is mentioned. A quarter

of the flyers in our sample are former peers of pilots who are mentioned at one point. Since pilots who are former peers of mentioned pilots are different from the rest of the sample (Appendix Table A.1), we always control for “ever peer” status or pilot fixed effects.

To obtain a first impression of a mention’s effect on peer pilots’ death rate, we estimate a Cox proportional hazard model:²¹

$$D_{it} = D_0(t)exp(\alpha_D E_i + \beta_D P_{it} + X_{it}\gamma_D) + \epsilon_{D,it} \quad (1)$$

Note that $D_0(t)$ stands for the baseline hazard function evaluated at month t (i.e. the baseline risk of death for any pilot t months after entering the war), α_D is a constant risk factor of ever having any one of your peers mentioned ($E_i = 1$), P_{it} is a dummy for a former peer of pilot i being mentioned in month t , and X_{it} is a vector of controls. Controls include dummies for the Eastern front, aircraft type, the pilot’s squadron, each month t of the war, and a measure of pilot quality, calculated as a pilot’s cumulative victories up to period $t-1$ divided by the number of months in combat.²² We also include a dummy for months with any pilot mention – which will absorb the effect from mentions changing an individual pilot’s relative standing in the air force *as a whole*. We are interested in the coefficient β_D , the extent to which a pilot’s death is more likely in a month when his former peer is publicly recognized.

To examine statistically the effect of a peer mention on the victory rate of pilot i in month t , we estimate the following least squares model:

$$V_{it} = \alpha_{V,i} + \beta_V P_{it} + X_{it}\gamma_V + \epsilon_{V,it} \quad (2)$$

Compared to the proportional hazard model, this equation includes also fixed effects for each individual pilot ($\alpha_{V,i}$), which implies that the control for ever peer status (E_i) is being absorbed. The remaining covariates X_{it} are the same, except that we also include a control for experience (the number of months a pilot has already been tracked in our data). Standard errors are clustered at the level of the squadron throughout all our specifications.²³

Table 1 presents our baseline results for the whole sample of pilots. Panel A shows results for exit rates. Months with (any) mention show somewhat higher exit rates, and pilots whose peers are mentioned at any one point in time exit at lower rates (columns 1 and 2).

²¹ Our estimates are virtually unchanged when we instead estimate a parametric survival-time model fitting a Weibull regression model using maximum likelihood. Results are available upon request.

²² Since death only happens once, the inclusion of pilot fixed effects for the risk of death would be nonsensical (and the Cox estimator does not converge). Instead, we control for ever peer status (E_i) when estimating the hazard model. We also do not control for a pilot’s experience since the model already takes into account the time at risk.

²³ Results are robust to two-way clustering (squadron and pilot), clustering by the identity of pilots mentioned in a month, clustering at the level of the base or to allowing for spatial autocorrelation (see Appendix Table A.2).

During the month of a former fellow pilot's mention, pilots see their hazard rate rise by more than 77%, on top of the general 16-22% rise in exit rates during mention periods (column 3). This effect becomes somewhat larger the more controls are added. In columns 6 and 7, we first include squadron fixed effects and then time fixed effects to our estimation. Even in the most demanding specification (column 7), when we control for pilot quality, front, aircraft type, and squadron and time fixed effects, we find significantly higher risks of exit (75% extra) for former peers of the mentioned pilot. The coefficient on past peer mentions suggests an increase in risk that is about three to five times larger than the one associated with general mention periods. This implies that intrinsic concerns about relative standing and personal rivalry (captured by the effect on past peers) may well be more powerful than other factors, such as a pilot's concern over his overall standing in the air force as a whole.

A similar pattern is visible for victory claims (Table 1, Panel B). Mention periods see more aerial victories in general (a quarter of a victory for the average pilot, col. 1). In months when a former peer is mentioned, the victory rate of his peers jumps by an additional half of a victory (column 2). After adding – in addition to individual fixed effects – controls for pilot quality (column 3) as well as experience, front, aircraft type, squadron and time fixed effects (column 6), having a former peer mentioned still adds more than a third of a victory in the same month. Again, we find that past peer effects are stronger than mention period effects, suggesting that intrinsic concerns over relative standing and personal rivalry are at least as important as instrumental factors.²⁴ Because both victory and exit rates rise after a peer's mention, high-powered incentives *may* have backfired: A full accounting of the overall efficiency effect would have to take into account the cost of training replacement pilots, their (time-varying) quality, and the aggregate impact of engineering a culture where status was closely tied to aerial victories. Neither parameter can be pinned down by our analysis.

B. Coarsened Exact Matching – Event Study

The results in Table 1 suggest large effects of former peer mentions on performance. However, they may reflect omitted variable, such as time- or state-dependent payoffs to unmeasured skills, which mentioned pilots and their peers share. To overcome such problems, we would ideally like to compare a pilot whose peer was mentioned with an identical, other pilot who happens not to have a mentioned peer.

To identify effects, for every treated pilot (i.e., a former peer of the mentioned pilot), we need to construct well-matched control groups. To this end, we first create *placebo mentions*. In every month, the Wehrmacht could have mentioned any active pilot in its

²⁴ Estimating with panel Poisson regressions yields similar results (see Appendix Table A.3). The same is true for estimating without the quality variable (Appendix Table A.4).

bulletins; only a few – if any – were actually mentioned. We first create placebo mentions from the set of plausible candidates, taking into account

1. the pilot’s average monthly victory rate up to time t
2. his number of victories in the current month
3. his number of victories in each of the preceding 3 months
4. the front where the pilot operated
5. the number of actually mentioned pilots in the same month

In the spirit of other papers creating comparison groups of high-performers (Azoulay et al. 2019; Jaravel et al. 2018), we coarsen these variables into a finite number of bins (using, for example, five bins for monthly victory scores).²⁵ Coarsened Exact Matching (CEM) was introduced by Iacus et al. (2012); in our setting, it ensures that there are enough placebo mention pilots for every actually mentioned pilot. To avoid contamination of the placebo and treated group, we restrict placebo-mentions to pilots who do not actually get mentioned (even at a later date).²⁶ We find matches for 72 out of 85 mention events. The attrition rate (13 out of 72, 15%) is lower than in other papers using CEM-matching.²⁷ To enhance the similarity between the groups, we use propensity score matching on the matches in each stratum, keeping the five most similar placebo-mentions (following Jäger 2016).²⁸ In the 72 actual treatment episodes with matches, 37 individual pilots were mentioned; we compare them to a group of 187 placebo-mentioned pilots, for whom we have 348 pilot-month observations during the placebo treatment.

Figures 1 and 2 illustrate our empirical strategy. Franz Doerr was mentioned in June 1944, having scored 20 victories in a month. Our procedure matches him to Ernst Weismann, whose time-path of performance looks similar in the months preceding his “placebo-mention” in August 1942. Walter Schuck had previously served with Doerr in the 7th squadron of JG 5, some 20 months earlier (Figure 1). In the months preceding Doerr’s mention, he is doing well, scoring 8, 11, and 10 victories per month. Then, when Doerr is mentioned, Schuck actually

²⁵ We coarsen the variables on victory rates in the current and three preceding months into five equal-width bins each. The variable on average victory rates up to the last month (our measure for pilot quality) is coarsened into five bins based on quality quintiles. The number of actually mentioned pilots is coarsened into a dummy that indicates whether at least one (other) pilot gets mentioned in this month. The front dummy is not coarsened.

²⁶ In this way, we can avoid coding months as placebo-mentions if performance during these months eventually led to an actual mention.

²⁷ For example, Azoulay et al. (2019) lose 25% of treated units due to a lack of suitable matches. Our matches come from 19 different strata.

²⁸ We estimate the propensity score with a Probit model and use uncoarsened versions of all our CEM variables as predictors. We keep slightly less than five times as many placebo-mentions than actual mentions, because a few strata have less than five potential matches.

scored 30 victories in a single month. Franz-Josef Beerenbrock previously flew with our placebo-mentioned pilot, Weismann; the two had served in 12th squadron, JG 51. Preceding the placebo mention of Weismann, Beerenbrock scored 5, 8, and 17 victories per month, very similar to Schuck's trajectory (Figure 2). Then, in the month when Weismann could have been mentioned, Beerenbrock's performance actually falls to 9 victories. Prior to the mention/placebo-mention, the performance of Beerenbrock and Schuck is statistically indistinguishable. During and after Doerr's mention, the former peer of the actually mentioned pilot (Schuck) outperforms the peer of the placebo-mentioned pilot (Beerenbrock) strongly. In essence, the empirical strategy for our event study repeats this exercise for 72 actually mentioned pilots, and plots average differences between actual peers and controls.

How similar are the mentioned pilots and the matched placebos? Figure 3, Panel A, plots one measure of ex-ante quality – pilot quality (defined as victories over months of flying experience prior to the mention episode) for actual vs placebo mentioned pilots. The grey boxes plot the range from the 25th to the 75th percentile; the outer “whiskers”, the 95% confidence interval. The distributions for actual and placebo-mentioned pilots overlap to a great extent. The same is true if we use a simpler measure, their victory score during the period of the mention (Panel B). When we examine balancedness, we find only small deviations (see Table 2). Placebo-mentioned pilots actually score *more* in the month of a potential mention than actually-mentioned pilots, an average of over 10 victories vs. 8 for the actually mentioned, during the month of the actual/potential mention ($t = 0$). If peers always reacted to spectacular performance of a fellow pilot, we would expect this to make it harder to find significant outperformance for the actual peers. During the run-up to the event, for each of the three preceding months, victory rates are statistically indistinguishable. Pilot quality is also similar, with the placebo-mentioned pilots actually having a slight edge. Both groups had similar exposure to the Eastern front, where victories were somewhat easier to achieve. Importantly, the frequency with which other, unconnected pilots are mentioned is also statistically indistinguishable.

We have thus created a group of placebo-mentioned pilots who closely resemble actually mentioned flyers. In a second step, we create a set of control pilots, reconstructing which pilots had, by the time of the placebo mention, flown with the placebo-mentioned pilots. We drop all pilots that are not either actual peers of a treated pilot, or part of the CEM-matched control group. To ensure a high degree of similarity, we also limit the control group

to pilots who at one point have flown with an actually mentioned peer.²⁹ While this approach does not create two treatment groups with identical characteristics, it greatly increases their similarity. We can compare treated and control pilots (a) when the actual mention occurs (treatment effect) and (b) when the placebo mention could have happened (placebo effect). Our procedure gives us 353 control and actual peer pilots, and more than 13,500 monthly observations. Note that we need to add “ghost pilots” to our dataset for the pre-period, i.e. fellow flyers of a placebo/actual mentioned pilot who were shot down prior to a mention, in order to analyze both victories and exits.

We can now analyze our two key outcomes in event time, restricting the sample to pilots who are either in the treated (actual peers of mentioned pilots) or control group (peers of placebo-treated pilots). To investigate the dynamics of the treatment effect, and to test if exits and victories of pilots with mentioned and placebo-mentioned peers followed parallel trends before the mention, we estimate an event-study specification:

$$y_{ite} = \sum_{k \neq -1} \beta_k \lambda_k \times P_{ie} + \lambda_t + \eta_i + \mu_e + \epsilon_{it} \quad (3)$$

The variable y_{ite} captures the outcomes of interest: how many victories pilot i scores, or whether he exits the sample, in month t of event e .³⁰ The months t are indexed in relative time to the event e of having a past peer receiving a mention or a placebo-mention. For example, observations of pilots one month before their former peer is mentioned are indexed $t = -1$. The dummy variable P_{ie} takes value 1 if the past peer of pilot i receives an actual mention in event e . If the event is for a placebo-mention, P_{ie} is equal to 0. Since we include individual pilot (η_i), event (μ_e) and event-time fixed effects (λ_t), the coefficients β_k will identify the effect of having a past peer actually receiving a mention on pilots' performance in month k . We leave out the coefficient for the period immediately preceding the mention, β_{-1} . This means that we always compare the difference in outcomes of pilots with mentioned and placebo-mentioned relative to the difference between them at $t = -1$. Standard errors are clustered at the level of the squadron.

Panel A of Figure 4 plots the set of β_k coefficients, i.e. the additional monthly exit rate of peers of an actually mentioned pilot, relative to the control group (peers of placebo-

²⁹ We drop a small number of control pilots who have a former peer getting actually mentioned in the same month that another former peer is placebo-mentioned. The following results are robust to keeping these pilots in the sample.

³⁰ Observations have to be indexed by event e and pilot i because some pilots experience multiple events in which their former peers is mentioned (or placebo-mentioned). This means that some pilot-months show up more than once in our sample. These additional observations do not lead us to understate our standard errors since we already cluster at a higher level than the pilot-month, the squadron.

mentioned pilots). There are no significant pre-trends. With the mention of an actual peer, the exit rate increases up to the third month, after which the effect starts to decline in magnitude. The coefficients are highly significant for months 0, 1, and 3 after the event, but become insignificant afterwards.³¹ For victories (Figure 4, Panel B), we similarly find a sharp change after an actual peer’s mention. There are no pre-trends. Victory rates jump after an actual mention, and the coefficients are positive for 6 months afterwards. For months 2 and 3, we find individually significant effects (at 10%).³²

Next, we estimate a difference-in-differences specification to examine the joint significance of effects of actual mentions compared with placebo-mentions, averaging effects across pre- and post-treatment periods. In particular, we estimate the following equation to recover the effect on performance of having an actual past peer mentioned:

$$y_{ite} = \beta P_{ie} \times post_t + \lambda_t + \eta_i + \mu_e + \epsilon_{it} \quad (4)$$

The dummy variable $post_t$ takes a value equal to 1 in months after a past peer of pilot i gets mentioned or placebo-mentioned. We are interested in the estimate of coefficient β of the interaction of $post_t$ and the dummy variable for treatment group status, P_{ie} . We again include individual pilot (η_i), event (μ_e), and event-time fixed effects (λ_t). With these fixed effects, β identifies the effect of mentions on the subsequent performance of former peers of the mentioned pilot.

Table 3 reports the difference-in-differences effects from Equation 4, confirming the results from our event-time plots. In Panel A, col. 1, we find that after a former peer gets mentioned, the risk of exit more than doubles (2.14 times). This is relative to the pre-mention difference in exit rates for peers of actual mentioned and placebo-mentioned pilots. The estimated effect is even larger (189% increase in exits) when we focus on the window of +/- 6 months around the mention event (col. 2).³³ Panel B of Table 3 reports our estimates for the effect on victory rates. Column 2 shows that for the six months after a pilot is honored with an

³¹ Our baseline results for exit rates are estimated with a Cox proportional hazard model. When estimating the Cox model, we deploy a more parsimonious version of Equation 3, where we replace the event and pilot fixed effects with a simple dummy for the treatment group, P_{ie} . The Cox model would not converge with the full set of fixed effects. Figure A.3 in the Appendix shows the equivalent plot when we estimate Equation 3 for exits with OLS. Since we are measuring survival, this is not our preferred specification but yields readily-interpretable coefficients. Each pilot can only exit once, and we therefore do not include P_{ie} instead of pilot fixed effect. The figure shows a sharp jump in the risk of exit in the treatment period and elevated effects for several months.

³² Appendix Figure A.4 replicates Figure 4 but clusters standard errors on the level of the mentioned (or placebo-mentioned) peer.

³³ As in Figure 4, Panel A, we also estimate Table 3, Panel A, with the Cox proportional hazard model and replace event and pilot fixed effects with a simple treatment group indicator. Appendix Table A.5 replicates the exit results with OLS estimates of Equation 4. Appendix Table A.6 replicates Table 3 but clusters on the level of the mentioned (or placebo-mentioned) peer.

actual mention, his former peers increase their monthly performance by an average 0.31 extra victories (compared with the performance of pilots whose former peer gets placebo-mentioned).³⁴

C. Birthplace proximity

We can use regional origin as a proxy for social proximity. While not every high-performing pilot knew every other top pilot, many of them would have been familiar with each other's careers and background. In addition, last names often convey information about regional origins. Here, we examine how much greater the increase in the number of victories is when a peer from the same region is mentioned.

We have geo-coded information on the birthplaces of 683 high-performing pilots. Figure 5 shows that for pilots born close to each other (i.e. less than 100 miles distance), the effect of a mention in dispatches on a former peer translates into an additional 0.5 victories per month.³⁵ Yet, at a distance of (say) 600 miles, the performance increase is smaller. Amongst the peers of placebo-mentioned pilots, we find no gradient in birthplace distance to the mentioned pilot. In fact, pilots do not react to mentions of individuals from the same region, unless the mentioned pilot is a former peer. Repeating the exercise with exits is possible, but not sensible in this setting. In our data set with birthplace data, only 7 pilots exit when their peer is mentioned; this does not allow us to quantify effects with any precision.

Honored individuals from one's region will return to the same social environment post-war – but those who are not (publicly recognized) peers do not motivate effort. This suggests that the underlying mechanism is not driven by social image concerns vis-à-vis others in one's region of birth (or by competition over instrumental benefits at the local level), but rather by *personal* rivalry.³⁶

D. Results by intensity of past interaction

So far, we have defined (former) peers exclusively as those who served together in the same squadron. Bonds between squadron peers were particularly close. At the same time, other forms of interaction may also have acquainted pilots with each other, possibly leading to bonding and/or rivalry.

³⁴ The results for the effect on victory rates are robust to focusing on a balanced panel of pilots around the +/-6 event window (see Appendix Table A.7). We cannot look at exits in the same way, since there are no exits in a balanced panel.

³⁵ We use the specification from Table 3, Panel B, column 1 because our sample otherwise becomes too small. We geo-coded birthplaces for 683 pilots, but only 217 of them remain in our sample after applying CEM. We can compute distances in birthplaces for 186 of them and the mentioned pilots.

³⁶ Of course, self-image concerns could be operative here – pilots may simply know more and care more about fellow pilots who are closer to them socially, even without reacting to changes in their standing in the eyes of others. Our finding is related to Bandiera et al. (2010), who show that in settings without production externalities, social ties between co-workers can strengthen peer effects.

We perform the same analysis as before, but for two other definitions of peers – pilots who previously served in the same group and those who flew from the same airbase. Groups consisted of 3-4 squadrons. They often flew together and would participate in joint training and recreational activities – but they would not necessarily fly from the same airfields (even if they often did so). Pilots from other groups would often use the same airbase, too, giving us another form of peer interaction. These would join in the same operation less frequently, but social interaction over a meal or a drink remained likely.

We again restrict ourselves to the CEM-matched sample. Figure 6 repeats the diff-in-diff analysis in Table 3, plotting the coefficient of interest for group peers and base peers, for both exit rates and victory rates. For exit rates, we find substantial effects for all groups. Squadron peers show the largest effect, followed by group peers and then, base peers. This is in line with our expectations – pilots who flew from the same base will have had some chances to interact, from drinking in the mess to joint outings; whereas group peers will have interacted frequently in training and in briefings, for example, plus shared the camaraderie of serving under the same officers and on related missions. Victory rates mirror the results for exits. We find increases for all types of peers, but those for squadron and group peers are largest – and the latter two are statistically indistinguishable from each other.

E. Results by pilot quality

Finally, we analyse our results by pilot quality. Some of the aces in our data were among the highest-scoring pilots of all time, who may have had substantial degrees of freedom to increase their score; others were junior pilots struggling to stay alive. We should not expect them to react in the same fashion.

In Figure 7, Panel A, we plot the exit rate of both groups in event time. Pilots in the bottom 80% show no significant pre-trend. After the mention of their peer, their exit rates jump markedly, by a factor of 6, and stay elevated for six months; four of the seven coefficients are highly significant. Top pilots show no pre-trend but in their case, there is no upward shift in exit rates after a mention. While six of the seven post-treatment coefficients are positive, none of them is significant.³⁷

For victories, we find the opposite pattern. Bottom 80% pilots show slight negative and insignificant deviations pre-treatment, and then a jump to several positive values, neither of which is significantly different from zero. For the top pilots, however, we see an increase in victories. While the treatment month itself and the next one are not significantly different

³⁷ Figure 7, Panel A, is estimated with the OLS specification. We replicate the same pattern of exits when estimating with Cox in Appendix Figure A.5.

from some of the coefficients pre-treatment, the victory rate jumps by more than 2 per month for months 2 and 3 relative to the last month pre-treatment.

Table 4 presents difference-in-differences estimates for the same sample, comparing pre- and post-mention performance. We find significant increases in exits for the bottom 80%, with a 166% rise, and significant changes in victory rates for both the top and bottom pilots – but much larger effects for leading ‘aces’.³⁸

IV. MECHANISMS, ADDITIONAL RESULTS, AND ALTERNATIVE INTERPRETATIONS

We next attempt to rule out potential confounding mechanisms. In addition, we examine the robustness of our findings.

A. Mechanisms

What margin of adjustment do pilots use to respond to mentions of former peers? Pilots could try harder on each flight they take; or they could fly more often/for longer. Distinguishing between risk-taking and effort is challenging – there is no risk-free way to try and down an enemy aircraft. Nonetheless, we offer some suggestive evidence to separate the two and find that both risk-taking and greater effort are behind the effects we find.

German air force pilots were required to log all their flights, including time of take-off and landing, number of sorties per day, type of mission, successes, and possible crashes. On any given day, pilots could fly three types of missions: combat flights, protection flights (cover for the air field), transfer flights (moving planes and equipment), and training flights. Victories were almost all scored on combat flights. For the majority of pilots, the flight logs have not survived; for those that did, they are almost exclusively held in private collections. We managed to purchase flight logs for a small subset of 71 pilots.

We find that the former peers of mentioned pilots increased their number of combat missions after a peer’s mention. In Figure 8, we plot the share of combat flights for the pilots in our sample, for 30 days before and after the mention of a peer. In total, we have 14 pilots with 428 daily observations of at least one flight in this interval.³⁹ Prior to the mention, about 50-60% of their flights are on combat missions; thereafter, this rises to more than 75% initially and remains elevated. Between the month of the mention and the previous month, the share of combat flights jumps by more than 15 percentage points.

³⁸ We replicate Figure 7 in Appendix Figure A.6, and Table 4 in Appendix Table A.8, but we cluster standard errors on the level of the mentioned (or placebo-mentioned) peer. Table A.9 in the Appendix shows alternative splits of our data – 90/10 and 70/30. Results remain largely unchanged.

³⁹ The remaining 57 pilots do not have past peers that get mentioned during the interval that we observe in their flight logs.

When we use the full set of 71 flight logs, we can examine how much the number of combat flights changes after a peer's mention. In Table A.10, we find increases that are (after standardization of the dependent variable) very similar in magnitude to changes in victory rates in our full panel. While the change is only statistically significant at 10% (p-value of 0.078), the coefficient is never significantly different from the increases in the full panel (p-value of 0.718). Since combat flights are highly correlated with victories overall, we conclude that increases in combat involvement are possibly one important contributor to the overall rise in victory rates after a former peer mention.

Fighter pilots had many margins of adjustment. They could volunteer for combat missions, or make judgment calls about the state of their aircraft, and decide to fly despite substantial damage to the airframe.⁴⁰ In our flight log sample (with N=7,886 daily observations of at least one flight), a regression of daily victories on the number of daily combat flights shows that they are strongly, positively correlated (with a t-statistic of 20.3; cf. Appendix Figure A.7).

Pilots flew more combat missions after a peer was mentioned, meaning that they put themselves more often in harm's way. Our main data allows us to confirm that the extensive margin of combat participation is a key driver of our results. While we do not observe flight frequency in our main data, we can calculate the number of active days, i.e. when a pilot claimed at least one victory. This can serve as a lower bound on the number of days with combat flights. When we compile this statistic, our CEM event-study framework shows the following pattern:

The number of days with at least one victory in a month increases strongly and discontinuously after a mention (Figure 9), rising by 0.2-0.4 extra days. It also stays elevated for at least 4 months after the month of the mention. This is in line with the flight log evidence, and suggests that the former peers of mentioned pilots exerted more effort by clocking more active days.

In addition, we can investigate the risk-taking channel. Every fighter pilot on a combat mission could take a number of decisions that increased his personal risk:

1. engage the enemy (at all)
2. try for victory even in an unpromising position, instead of breaking away

⁴⁰ Heinz Knoke landed his plane with a broken spar in 1943. Nonetheless, he decided to take off when he saw American bombers overhead. He recalled: "It makes me sick to think of my plane as unserviceable. Suddenly, I make up my mind that I am going to fly in my damaged plane nevertheless. Despite the warnings of the inspector, I have it refuelled and rearmed." (Knoke 1953). He then took off to attack American B-17 bombers, destroying one before crash-landing his plane. Conversely, pilots seeking to avoid combat would often claim engine trouble or problems with their oxygen (McFarland and Newton 2006).

3. try to score multiple victories, instead of heading home after the first (pilots are often at the limit of physical endurance even after a single engagement)
4. continue to attack another aircraft even after substantial damage to their own plane

These decisions are not directly observed. If a pilot takes greater risks in any or all of these dimensions, victory rates may go up, but so will the risk of being shot down. We compile a new indicator, at monthly frequency, of the ratio of exits/days with victories. This captures whether, conditional on flying a combat mission, a pilot experienced a higher risk of being shot down. We perform the analysis in the same CEM framework as the main results of the paper, limiting the data to the treated and control groups. We find that the risk of exit per combat day jumps discontinuously by over 5 percentage points immediately afterwards, and stays elevated in subsequent months (Figure 10).⁴¹

In combination, these two results suggest that risk-taking increased substantially after a peer's mention. These findings add to the evidence from the flight logs documenting greater effort. Overall, we find substantial evidence to conclude that both margins – greater effort, and more risk-taking – led to increased victory rates and higher exits after the public mention of a former peer.

B. Rivalry vs social image concerns

Pilots react to events that change their position relative to a former peer whom they know well. This could either reflect a pilot's concern about how he is perceived by his "social universe" – or an intensely personal, private reaction to a decline in relative standing. To help investigate which of these two is responsible, we track the size and composition of each pilot's social universe – the total number of pilots he ever flew with, and the number of *other* pilots who are/were also peers of the mentioned pilot. We then examine whether the size of the effects we find is affected by the share of each pilot's social universe that also served with the honored pilot. If the relative *social* standing channel is key, we would expect an interaction of a peer being mentioned and that peer's "social footprint" in a pilot's universe of fellow pilots to be positive – pilots should react more to a former peer's mention if most of the flyers they know are also peers. Conversely, if this coefficient is zero or negative, the most likely interpretation is an intensely personal, private reaction to rivalry. While the size of a pilot's social universe is not randomly assigned, this approach may shed further light on the proposed mechanism.

⁴¹ Note that this is not driven by more victories per combat day, which remain broadly stable – scaling the number of exits by the number of victories, we obtain a very similar pattern. The ratio of exits to victories jumps by over 5 percentage points immediately after a mention (Appendix Figure A.8).

Figure 11 gives the results of adding the interaction effect to our baseline specification. We find that the higher the share of fellow pilots that also served with a mentioned pilot, the *lower* a pilot's exit rate and performance bump during a peer's mention is.⁴² In the top panel, we show that the share of pilots who also knew the mentioned pilot reduces the effect of mentions on exit rates, by up to 2/3rd. The effect of mentions on exit rates is significantly greater than zero if pilots share less than 68 percent of their social universe with the mentioned peer. This is the case for 72 percent of our sample. A similar result is obtained for victory rates (Figure 11, bottom panel). Here, we see a decline by up to 0.9 aircraft per month shot down for pilots whose entire social universe is composed of former peers of the mentioned pilot. To put this into perspective, an increase in shared peers from zero to 50% decreases the treatment effect from 0.87 to 0.40 victories. Although we cannot fully rule out a role for social image concerns, these results strengthen the case for the alternative interpretation – “rivalry,” or a private, personal reaction, as a key driver of risk-taking and effort.

D. Correlated shocks

Unobserved, correlated shocks could simultaneously affect the outcomes of different peer groups. While we exclude pilots serving in the same squadron when looking at past peers, this may not be enough to rule out the effect of aggregate changes in the combat environment.

To see if correlated shocks might drive our findings, we can exclude pilots from nearby units. For this purpose, we impose a minimum distance requirement between airfields from which the treated pilot and his former peers operated. During World War II, German forces were fighting from the Arctic Circle to the deserts of North Africa and from Stalingrad to the Pyrenees. Figure 12 shows the location of Luftwaffe airfields on a map of Europe and North Africa. The minimum distance between airfields in our data is less than a mile, and the maximum is 2,753 miles.

Figure 13 examines peer interactions as a function of minimum distance requirements between the airfields of two (former) peers, using our CEM sample. Even a distance of 100 miles usually corresponded to a marked change in combat conditions. At a distance of 500 miles, units would be operating with different army groups (North, Center, or South) on the Eastern front. Units flying bomber intercept missions over Germany were separated by up to 1,000 miles from their counterparts on the Eastern front. The effects on exits shown in Figure 13, Panel A, are similar across distance groups (and not statistically different from each

⁴² A pilot's “social universe” is correlated with other important characteristics. However, both exits and victory rates are *positively* correlated with the size of a pilot's social universe overall. Since we are estimating with pilot fixed effects (for victories), results are not reflective of unbalanced sample composition.

other). Panel B on victories shows no clear pattern for the size of the coefficient; no sub-sample result is significantly different from our baseline estimate. These results strongly suggest that our results are not driven by correlated shocks.

Aircraft changes could also affect our results. Performance could reflect improvements in technology. Thus, a sudden increase in the number of aerial victories could be driven by good pilots receiving simultaneous upgrades in their planes. This, however, is unlikely to explain our results. We have information on the aircraft type for 77,000 of our total 82,000 monthly observations (see Appendix Figure A.9 for the distribution of aircraft types used). Most missions were flown in one of four types – the BF-109E, F, and G and the FW-190. The Luftwaffe typically upgraded entire squadrons to facilitate maintenance and training. There is no anecdotal evidence of top performing pilots being given special treatment.⁴³

V. CONCLUSION

Using new data on confirmed aerial victories and losses by German pilots during World War II, we study how personal rivalry can motivate both effort and risk-taking in a high-stakes setting. Our results show that pilots respond strongly to the public recognition of peers they are personally familiar with. A mention in the daily bulletin of the German armed forces for outstanding accomplishments was considered a great honor. Pilots who previously flew with the mentioned pilot begin to score more victories immediately afterwards. At the same time, their risk of death increases considerably. These effects vary by skill group: Performance gains are concentrated among highly skilled pilots; and while average pilots also score more, their gains are relatively small. Risk increased significantly for the low-skilled pilots. In other words, inspired by the accomplishments of German air force ‘aces’ to try harder, average pilots won few additional victories, but perished at a much higher rate.

Tangible benefits linked to relative standing are not influenced by a recognized pilot’s personal connection with other pilots. Our empirical strategy demonstrates the importance of *intrinsic* positional concerns and *personal* rivalry – pilots try much harder to restore their relative standing vis-à-vis others who they know personally, compared with anonymous competitors. The greater the similarity in background, the greater the increase in risk-taking and effort: Pilots react more the closer their birthplace is to that of the mentioned pilot, even after controlling for mentions of other pilots. Effects are not amplified if a pilot’s social

⁴³ At least one leading pilot (Hans-Joachim Marseille) was, despite his protests, forced to pilot a new BF-109G because his entire squadron was being re-equipped. Marseille died shortly thereafter when the new plane’s engine failed. In Table 1, we explicitly control for plane FE.

universe largely consists of a mentioned pilot's former peers – the reaction is more likely to reflect *personal rivalry* than relative social concerns within a group.

Personal rivalry is an everyday occurrence – from teenage pranks to professional rivalries among hedge fund traders, lawyers, doctors, and researchers. Our findings demonstrate that *private, intrinsic* relative rivalry can be a major source of motivation for highly skilled individuals. Importantly, it may also increase risk-taking. Our particular setting, with detailed information on high frequency outcomes, allows us to examine how far people are prepared to go to improve their relative standing. Our data suggest that – in a sample of highly motivated pilots – many are prepared to risk their life to enhance their relative position vis-a-vis a personally known “rival”. This effect also underlines the potential for high-powered incentives to backfire in settings where both performance and risk are important for aggregate effectiveness⁴⁴ – personal rivalries may incentivize individuals too strongly, leading to a (potential) decline in aggregate performance.

References

- Acemoglu, Daron, Leopoldo Fergusson, James Robinson, Dario Romero, and Juan F. Vargas. 2020. “The Perils of High-Powered Incentives: Evidence from Colombia’s False Positives.” *American Economic Journal: Economic Policy* 12 (3): 1–43.
- Agarwal, Sumit, Vyacheslav Mikhed, and Barry Scholnick. 2016. “Does Inequality Cause Financial Distress? Evidence from Lottery Winners and Neighboring Bankruptcies.” *Philadelphia Fed Working Paper* 16–04: 1–47.
- Andrew, Dean Wing Commander. 2004. “Strategic Culture in the Luftwaffe – Did It Exist in World War II and Has It Transitioned into the Air Force?” *Defence Studies* 4 (3): 361–86.
- Ashraf, Nava, Oriana Bandiera, and B. Kelsey Jack. 2014. “No Margin, No Mission? A Field Experiment on Incentives for Public Service Delivery.” *Journal of Public Economics* 120: 1–17.
- Azoulay, Pierre, Christian Fons-Rosen, and Joshua S. Graff Zivin. 2019. “Does Science Advance One Funeral at a Time?” *American Economic Review* 109 (8): 2889–2920.
- Azoulay, Pierre, Jialan Wang, and Joshua S. Graff Zivin. 2010. “Superstar Extinction.” *The Quarterly Journal of Economics* 125 (2): 549–89.
- Bandiera, Oriana, Iwan Barankay, and Imran Rasul. 2005. “Social Preferences and the Response to Incentives: Evidence from Personnel Data.” *The Quarterly Journal of Economics* 120 (3): 917–62.
- . 2010. “Social Incentives in the Workplace.” *Review of Economic Studies* 77 (2): 417–58.
- Bertrand, Marianne, and Adair Morse. 2016. “Trickle-Down Consumption.” *The Review of Economics and Statistics* 98 (5): 863–79.
- Besley, Timothy, and Maitreesh Ghatak. 2008. “Status Incentives.” *American Economic Review* 98 (2): 206–11.

⁴⁴ One related example is competition amongst traders as in Kirchler et al. (2018).

- Bottan, Nicolas L., and Ricardo Perez-Truglia. 2020. "Choosing Your Pond: Location Choices and Relative Income." *The Review of Economics and Statistics*, October, 1–46.
- Brown, Jennifer. 2011. "Quitters Never Win: The (Adverse) Incentive Effects of Competing with Superstars." *Journal of Political Economy* 119 (5): 982–1013.
- Brown, Keith C., W.V. Harlow, and Laura T. Starks. 1996. "Of Tournaments and Temptations: An Analysis of Managerial Incentives in the Mutual Fund Industry." *Journal of Finance* 51 (1): 85–110.
- Bursztyn, Leonardo, Bruno Ferman, Stefano Fiorin, Martin Kanz, and Gautam Rao. 2018. "Status Goods: Experimental Evidence from Platinum Credit Cards." *The Quarterly Journal of Economics* 133 (3): 1561–95.
- Bursztyn, Leonardo, and Robert Jensen. 2015. "How Does Peer Pressure Affect Educational Investments." *The Quarterly Journal of Economics* 130 (3): 1329–67.
- Caldwell, Donald L. 1996. *The JG 26 War Diary*. London: Grub Street.
- . 2012. *Day Fighters in Defence of the Reich: A War Diary, 1942–45*. Barnsley, South Yorkshire: Frontline Books.
- Card, David, Alexandre Mas, Enrico Moretti, and Emmanuel Saez. 2012. "Inequality at Work: The Effect of Peer Salaries on Job Satisfaction." *American Economic Review* 102 (6): 2981–3003.
- Chan, Ho Fai, Bruno S. Frey, Jana Gallus, and Benno Torgler. 2014. "Academic Honors and Performance." *Labour Economics* 31 (December): 188–204.
- Costa, Dora L., and Matthew. E. Kahn. 2003. "Cowards and Heroes: Group Loyalty in the American Civil War." *The Quarterly Journal of Economics* 118 (2): 519–48.
- DellaVigna, Stefano, John A. List, and Ulrike Malmendier. 2012. "Testing for Altruism and Social Pressure in Charitable Giving." *The Quarterly Journal of Economics* 127 (1): 1–56.
- . 2017. "Voting to Tell Others." *Review of Economic Studies* 84: 143–81.
- deZeng, Henry L. IV, and Douglas G. Stankey. 2019. "Luftwaffe Officer Career Summaries." April 1, 2019. <http://www.wv2.dk/lwoffz.html>.
- Duesenberry, James S. 1949. *Income, Saving and the Theory of Consumer Behavior*. Cambridge, MA: Harvard University Press.
- Ederer, Florian. 2010. "Feedback and Motivation in Dynamic Tournaments." *Journal of Economics & Management Strategy* 19 (3): 733–69.
- Enikolopov, Ruben, Alexey Makarin, and Maria Petrova. 2020. "Social Media and Protest Participation: Evidence From Russia." *Econometrica* 88 (4): 1479–1514.
- Evans, Richard. 2009. *The Third Reich at War: 1939-1945*. New York, NY: Penguin.
- Falk, Armin, and Andrea Ichino. 2006. "Clean Evidence on Peer Effects." *Journal of Labor Economics* 24 (1): 39–57.
- Fershtman, Chaim, and Uri Gneezy. 2011. "The Tradeoff Between Performance and Quitting in High Power Tournaments." *Journal of the European Economic Association* 9 (2): 318–36.
- Festinger, Leon. 1954. "A Theory of Social Comparison Processes." *Human Relations* 7 (2): 117–40.
- Frey, Bruno S. 2007. "Awards as Compensation." *European Management Review* 4 (1): 6–14.
- Galland, Adolf. 1993. *Die Ersten und die Letzten: die Jagdflieger im Zweiten Weltkrieg*. München: Schneekluth.
- Genakos, Christos, and Mario Pagliero. 2012. "Interim Rank, Risk Taking, and Performance in Dynamic Tournaments." *Journal of Political Economy* 120 (4): 782–813.
- Goltsman, Maria, and Arijit Mukherjee. 2011. "Interim Performance Feedback in Multistage Tournaments: The Optimality of Partial Disclosure." *Journal of Labor Economics* 29 (2): 229–65.

- Green, Jerry R., and Nancy L. Stokey. 1983. "A Comparison of Tournaments and Contracts." *Journal of Political Economy* 91 (3): 349–64.
- Iacus, Stefano M., Gary King, and Giuseppe Porro. 2012. "Causal Inference without Balance Checking: Coarsened Exact Matching." *Political Analysis* 20 (1): 1–24.
- Jackson, C. Kirabo, and Elias Bruegmann. 2009. "Teaching Students and Teaching Each Other: The Importance of Peer Learning for Teachers." *American Economic Journal: Applied Economics* 1 (4): 85–108.
- Jäger, Simon. 2016. "How Substitutable Are Workers? Evidence from Worker Deaths." *SSRN Electronic Journal*. <http://www.ssrn.com/abstract=2925897>.
- Jaravel, Xavier, Neviana Petkova, and Alex Bell. 2018. "Team-Specific Capital and Innovation." *American Economic Review* 108 (4–5): 1034–73.
- Karaca-Mandic, Pinar, Nicole Maestas, and David Powell. 2013. "Peer Groups and Employment Outcomes: Evidence Based on Conditional Random Assignment in the U.S. Army." *Working Paper*.
- Kirchler, Michael, Florian Lindner, and Utz Weitzel. 2018. "Rankings and Risk-Taking in the Finance Industry." *The Journal of Finance* 73 (5): 2271–2302.
- Knoke, Heinz. 1953. *I Flew for the Führer: A Story of a German Airman*. Longon: Evans.
- Knowles, Elizabeth, ed. 1999. *The Oxford Dictionary of Quotations*. 5th ed. Oxford; New York: Oxford University Press.
- Kosfeld, Michael, and Susanne Neckermann. 2011. "Getting More Work for Nothing? Symbolic Awards and Worker Performance." *American Economic Journal: Microeconomics* 3 (3): 86–99.
- Kracker Luftwaffe Archive, last modified Januar 22, 2016. n.d. <http://www.aircrewremembered.com/KrackerDatabase/>.
- Kroener, Bernhard, Rolf-Dieter Müller, and Hans Umbreit. 1988. *Organisation Und Mobilisierung Des Deutschen Machtbereichs*. Das Deutsche Reich Und Der Zweite Weltkrieg, Bd. 5. Stuttgart: Deutsche Verlags-Anstalt.
- Kuhn, Peter, Peter Kooreman, Adriaan Soetevent, and Arie Kapteyn. 2011. "The Effects of Lottery Prizes on Winners and Their Neighbors: Evidence from the Dutch Postcode Lottery." *American Economic Review* 101 (5): 2226–47.
- Kurowski, Franz. 1994. *German Fighter Ace: Hans-Joachim Marseille: Star of Africa*. Atglen, PA: Schiffer.
- Kuziemko, Ilyana, Ryan W. Buell, Taly Reich, and Michael I. Norton. 2014. "'Last-Place Aversion': Evidence and Redistributive Implications." *The Quarterly Journal of Economics* 129 (1): 105–49.
- Lazear, Edward P., and Sherwin Rosen. 1981. "Rank-Order Tournaments as Optimum Labor Contracts." *Journal of Political Economy* 89 (5): 841–64.
- Lizzeri, Alessandro, Margaret A. Meyer, and Nicola Persico. 2002. "The Incentive Effects of Interim Performance Evaluations." *Penn CARESS Working Paper* 02–09: 1–25.
- Luttmer, Erzo. F. P. 2005. "Neighbors as Negatives: Relative Earnings and Well-Being." *The Quarterly Journal of Economics* 120 (3): 963–1002.
- Malmendier, Ulrike, and Geoffrey Tate. 2009. "Superstar CEOs." *The Quarterly Journal of Economics* 124 (4): 1593–1638.
- Markham, Steven E., K. Dow Scott, and Gail H. McKee. 2002. "Recognizing Good Attendance: A Longitudinal, Quasi-Experimental Field Study." *Personnel Psychology* 55 (3): 639–60.
- Mas, Alexandre, and Enrico Moretti. 2009. "Peers at Work." *American Economic Review* 99 (1): 112–45.
- Mathews, Johannes, and John Foreman. 2015. *Luftwaffe Aces*. Leicester, UK: Wing Leader.

- McFarland, Stephen L, and Wesley Phillips Newton. 2006. *To Command the Sky The Battle for Air Superiority Over Germany, 1942-1944*. Tuscaloosa, AL: University of Alabama Press.
- McPherson, James M. 1997. *For Cause and Comrades: Why Men Fought in the Civil War*. New York, NY: Oxford University Press.
- Moldovanu, Benny, Aner Sela, and Xianwen Shi. 2007. "Contests for Status." *Journal of Political Economy* 115 (2): 338–63.
- Murray, Williamson. 1996. *The Luftwaffe, 1933-45: Strategy for Defeat*. Potomac Books.
- Nalebuff, Barry J., and Joseph E. Stiglitz. 1983a. "Information, Competition, and Markets." *American Economic Review* 73 (2): 278–83.
- . 1983b. "Prizes and Incentives: Towards a General Theory of Compensation and Competition." *The Bell Journal of Economics* 14 (1): 21.
- Neckermann, Susanne, Reto Cueni, and Bruno S. Frey. 2014. "Awards at Work." *Labour Economics* 31 (December): 205–17.
- Perez-Truglia, Ricardo. 2020. "The Effects of Income Transparency on Well-Being: Evidence from a Natural Experiment." *American Economic Review* 110 (4): 1019–54.
- Perez-Truglia, Ricardo, and Guillermo Cruces. 2017. "Partisan Interactions: Evidence from a Field Experiment in the United States." *Journal of Political Economy* 125 (4): 1208–43.
- Salter, James. 1956. *The Hunters*. 1st Edition. Harper & Brothers.
- Scherzer, Veit. 2005. *Ritterkreuzträger 1939-1945*. Ranis: Scherzers Militaer-Verlag.
- Shue, Kelly. 2013. "Executive Networks and Firm Policies: Evidence from the Random Assignment of MBA Peers." *The Review of Financial Studies* 26 (6): 1401–42.
- Simkin, Mikhail V., and Vwani P. Roychowdhury. 2006. "Theory of Aces: Fame by Chance or Merit?" *The Journal of Mathematical Sociology* 30 (1): 33–42.
- Smith, Adam. 1759. *The Theory of Moral Sentiments*. London: Printed for A. Millar and A. Kincaid and J. Bell.
- Spick, Mike. 2011. *Luftwaffe Fighter Aces: The Jagdflieger and Their Combat Tactics and Techniques*. Barnsley, South Yorkshire: Frontline Books.
- Stedman, Robert F., and Mike Chappell. 2002. *Luftwaffe Air & Ground Crew 1939-45*. Men-at-Arms 377. Oxford: Osprey Publishing.
- Stouffer, Samuel A., Arthur A. Lumsdaine, Marion H. Lumsdaine, Robin M. Williams Jr, M. Brewster Smith, Irving L. Janis, Shirley A. Star, and Leonard S. Cottrell Jr. 1949. *The American Soldier: Adjustment During Army Life*. Vol. 1. New Jersey, NJ: Princeton University Press.
- USSBS. 1945. *The United States Strategic Bombing Survey*. Vol. 56. Washington: US Government Printing Office.
- Van Creveld, Martin. 2007. *Fighting Power: German and US Army Performance, 1939-1945*. Contributions in Military History, no. 32. Westport, CT: Greenwood Press.
- Veblen, Thorstein. 1899. *The Theory of the Leisure Class: An Economic Study in the Evolution of Institutions*. New York, NY: Macmillan.
- Waldinger, Fabian. 2012. "Peer Effects in Science: Evidence from the Dismissal of Scientists in Nazi Germany." *Review of Economic Studies* 79 (2): 838–61.
- Wegmann, Günter, ed. 1982. "Das Oberkommando Der Wehrmacht Gibt Bekannt": *Der Deutsche Wehrmachtbericht: Vollständige Ausgabe Der 1939-1945 Durch Presse Und Rundfunk Veröffentlichten Texte*. Osnabrück: Biblio Verlag.
- Wood's Combat Claims, last modified October 21, 2013. n.d. <https://web.archive.org/web/20130928070316/http://lesbutler.co.uk/claims/tonywood.htm>.
- Yildirim, Huseyin. 2005. "Contests with Multiple Rounds." *Games and Economic Behavior* 51 (1): 213–27.