9 November 2026

Methodological Issues in Resolving Deer, Wolf, and People Avoidance from Camera Records: from One Home Ground to County and then a Statewide Volunteer Network

Damian A. Vraniak

Research Institute on Sustainable Ecological Relationships,

Vraniak Prairies & Savannas, W3177 Hamilton Road, Springbrook, WI 54875, USA

Manuscript submitted to Wildlife Society Bulletin; also deposited as a preprint on EcoEvoRxiv (doi to follow).

Corresponding author: Damian A. Vraniak, Research Institute on Sustainable Ecological Relationships, Vraniak Prairies & Savannas, W3177 Hamilton Road, Springbrook, WI 54875, USA. mashkiki75@gmail.com

Word count (main text, after cuts): 8,721

ABSTRACT Camera records are increasingly used to measure how wildlife respond to people and to predators, but the measures are sensitive to choices that are rarely stated: these include what counts as one event, what the delay is divided by, to which species a return-time measure can justifiably be applied, and against what null a directional test is compared. This paper reports those choices as informed by one dataset. The dataset was extended outward in four nested steps — a family's home ground (ten cameras, 2024–2026), one county (71 cameras), five counties (525 cameras) and the statewide wolf range (5,371 cameras) of the Snapshot Wisconsin

volunteer network, 2018–2025. The numbers from each wider extract were required to reproduce the narrower one's published values before any new question was asked. Seventeen issues are reported with the result each one changed; eight of them appear in no deposited paper. A checklist for volunteer-network camera analyses closes the paper.

KEYWORDS activity overlap, camera trap, *Canis lupus*, gray wolf, human disturbance, *Odocoileus virginianus*, pseudoreplication, Snapshot Wisconsin, volunteer networks, white-tailed deer.

1. Introduction

Automatically triggered cameras first used in wildlife studies supported making inferences in relation to inventories of the species present at a site and where animals could be told apart as individuals, estimating abundance by capture–recapture (Karanth 1995). Using cameras for detection probability entered later, once it was recognized that a species not photographed is not thereby absent (MacKenzie et al. 2002; Burton et al. 2015). A third use has increased with employing volunteer networks focusing on behavioral interactions within and between species, as well as between animals and people. Within that third use, three kinds of measure are relevant to what is reported here. The first is activity and overlap: the shape of a species' day, and the share of it that coincides with another species' day (Ridout and Linkie 2009; Rowcliffe et al. 2014; Frey et al. 2017; Lashley et al. 2018). The second is the detection rate: how often a species appears per unit of effort, and what that frequency can and cannot stand in for, relative abundance included (Sollmann et al. 2013; Burton et al. 2015; Sollmann 2018; Kays et al. 2020). The third is the passage-keyed response: what an animal does after something else has gone by

— return later, return at a different hour, or not return (Parsons et al. 2016; Gaynor et al. 2018; Naidoo and Burton 2020; Gilbert et al. 2022).

Each kind of measure rests on decisions taken before doing an analysis - what counts as one event, what the delay is divided by, which species the measure can honestly be applied to, and what null a directional claim is compared against. These decisions are usually made once, early, by whoever writes the script, and they are usually not reported. Several of the checks reported here move a P-value across the conventional threshold, and two of them reversed the interpretation of a result.

Volunteer networks across an entire state produce a dataset large enough that almost any contrast can be made to appear somewhere in it, and large enough that no reader can check the results obtained by simply looking at the photographs (Iannarilli et al. 2021; Townsend et al. 2021).

This paper reports the choices made exploring one dataset, extended outward in four nested steps: a family's own home ground, every frame labeled by the people who took it; the county that contains it; the five counties around it; and the statewide wolf range — the same extraction script with the county filter removed. Table 2 describes the data each extract held. The findings are located in five deposited papers (Vraniak 2026a–e) and in the home-ground camera census (Vraniak, 2026f).

Five terms define the measurements in this paper, and the same words apply to every species, including people. A *photograph* (or frame) is one stored image, labeled to a species, a camera and a time. A *visit* is an independent set of photographs of one species at one camera: a photograph begins a new visit only if it falls thirty minutes or more after the first photograph of the previous visit of that species at that camera, and a photograph of any other species in between

neither ends nor restarts it (§1.1 states the rule, its anchor and the alternatives; §3.16 details the cost). An independent detection or an event is a visit under another name; those terms are used in what follows only where a deposit's value is quoted in its own terms. A *pass* is the one unit that is not camera-level: on the home ground a person or a machine on a circuit sweeps every station in one trip, and that trip's photographs across all stations are collapsed into one pass (§3.2). A *dataset* is the whole of the photographs held at one scale: the home ground, Washburn County, the five counties, or the statewide wolf range. The three wider datasets are also called *extracts*, because each was drawn from the Snapshot Wisconsin network by the same script.

This paper's contribution to camera-trap methodology is the set of choices I made in examining one dataset that expanded outward from a family's home ground, and the checks I then ran against those choices, each stated with the numbers it changed, in the order a widening dataset changed them. Eight of the following checks appear in no deposited paper: 1) treating each human pass through the system, rather than each photograph or each camera, as the unit of independent replication (§3.2); 2) reporting every response index (the delay before an animal returns after a person or another species has passed, divided by that animal's usual interval at the same camera (§3.1) beside the coefficient of overlap between the two species' daily activity curves, so that a delay is not read where the two rarely keep the same hours (§3.4); 3) checking a small, verified local dataset against the statewide by-hour and by-month distributions published by Snapshot Wisconsin before the two are compared (§3.6); 4) asking, with uncertainty clustered by camera, which directional results hold across many cameras and which rest on a few (§3.8); 5) repeating every result resolved by sex or deer class on a neighboring seasonal window, to see whether it depends on where the season's boundary was drawn (§3.13); 6) clustering cameras at the spatial scale over which the landscape covariates are measured, nine square miles,

rather than treating each camera as independent (§3.15); 7) keying each event within one species, and showing that this choice cancels exactly from the index when it is read against the wild-turkey reference (§3.16); and 8) stating, for each landscape covariate, both the area it is measured over and the area of land it is meant to represent (§3.17).

The paper is organized so that each numbered issue in §3 is one section of the results, in the order the data raised the question, and each carries four parts: the problem as a reader meets it; what was done; the numbers it changed, cited to the paper that holds it; and the lesson, stated once. A checklist (§4) presents the lessons in the order an analyst would meet them. Three of the seventeen — §§3.15 to 3.17 — were not raised by the data but by the first reading by a reviewer, Adrian Treves, of this work from outside it, in September 2026, and presented at the end of §3 for that reason.

1.1 What independence means here, and why the rest depends on it

In part, the challenge of pseudoreplication (Hurlbert 1984) is telling when two photographs are of the same thing happening and when they are of two separate things. A camera fires in bursts; a wolf that walks past may leave one photograph or nine; a family of deer crossing a trail may leave forty. Counted as they occur, a busy minute becomes a busy season, and any test run on those counts reports the camera's trigger settings as though they were the animals' behavior.

The usual repair is a rule: photographs of the same species at the same camera are collapsed into one event unless they are separated by some interval, conventionally thirty minutes (O’Brien et al. 2003; but see Treves et al. 2010, who show that gregarious species inflate their counts under it, and who kept photographs of one species at one station on different days as the independent unit). The rule is a convention that can be set empirically, from the distribution of intervals

between a species' own photographs at the site (O. Deering, University of Wisconsin–Madison, unpublished data, set the interval at 21 minutes from the distribution of 433 consecutive-photograph intervals in the first months of a wolf-pack study on private land in southern Wisconsin; A. Seaborg, unpublished data, carried the same interval to the residential parcels around that pack's core). There are three further questions underneath it that no single number answers - whether the interval is measured from the last photograph or from the last one kept, whether a photograph of some other animal in between ends the event, and whether the cameras themselves — hosted by volunteers a mile or two apart, and each known only to the square mile — are separate observations or several views of the same piece of country (§3.15). Sections 3.15 to 3.17 take those three in turn.

To answer the question 'how many times did a wolf come?', suppose a camera records a wolf at 8:02, a wolf at 8:09, a deer at 8:20, a wolf at 8:24, a wolf at 8:41, a person at 9:40, a wolf at 10:05 and a wolf at 11:15. The question can be answered this way:

Rule for counting wolf events	Wolf events	What that rule is counting
Every photograph is an event	6	the camera's trigger rate, not the wolf's visits
15 minutes since the last kept photograph	5	8:09 falls inside the interval; 8:24 and 8:41 do not
30 minutes since the last kept photograph (the rule used here)	4	8:09 and 8:24 fall inside; 8:41 is 39 minutes after 8:02
30 minutes since the previous photograph, kept or not	3	the burst rolls the clock forward and swallows 8:41
30 minutes, but a photograph of any other species ends the event	5	the deer at 8:20 makes 8:24 a new wolf event; the wolf's count now rises with the deer's traffic

Rule for counting wolf events	Wolf events	What that rule is counting
60 minutes since the last kept photograph	3	the whole morning cluster becomes one event

Table 1. Four of the six numbers describe the same morning. The measures in this paper use the third rule, for three reasons: it is the rule of the deposited papers, so their values reproduce; it runs from the last kept photograph, so a long bout at one camera is broken into periodic events rather than swallowed into one (row four); and it is keyed within species, so the wolf's count does not rise with the deer's traffic (row five). Section 3.16 reports what the other choices do to the published figures, and why one class of figure is untouched by all of them.

2. The datasets at four scales

The four extracts are nested, not independent, each issue raised at one scale and tested at another: the home ground sits inside Washburn County, which sits inside the five-county block, which sits inside the statewide wolf range. Each widening exploration asked the same questions of a larger and more loosely supervised dataset. The narrower, home dataset — where the labels were made by the people who live there and can be checked against the photographs uploaded by them— served as the calibration for the wider ones.

Table 2. The four extracts, what each held, and the papers that used them.

	Home ground	Washburn County	Five counties (Barron, Burnett, Bayfield, Douglas, Sawyer)	Statewide wolf range
Source	Family's own cameras; every frame labeled	WI DNR Office of Applied Science, Snapshot Wisconsin extract	Same, wolf-enriched extract	Same extraction script, county filter removed
Cameras · effort	~10 cameras, 9 stations; 127 camera-model-months	71 cameras	525 cameras; 294,031 camera-days	5,371 cameras
Detections	11,820 frames; 5,499 independent deer detections	73,132 detections	634,214 detections	7,791,771 detections (190,596 coyote; 33,243 bear; 18,516 bobcat; 452,275 turkey)
Wolves	2 independent detections (4 frames)	55 detections; 2 sequential wolf–deer events	1,856 detections on 197 cameras; 116 sequential encounters	7,742 detections on 810 cameras; 461 sequential encounters
People	1,027 frames; 670 independent passes	7,180 human detections	432 cameras with a human-return ratio	2,847 cameras with a turkey–human ratio; deer–human ratio 1.82–1.92×
Years	2024 – Jun 2026	2018–2025	2018–2025	2018–2025
Used in	Vraniak 2026a; 2026f	Vraniak 2026b (county rung)	Vraniak 2026b, 2026c	Vraniak 2026d, 2026e

It is important to determine what remains true from one scale to the next, and what does not. Nothing measured on the home ground was carried to another camera as a value. Each index in

this paper was computed afresh at every camera at every scale, against that camera's own baseline (§3.1, §3.10), and only then summarized across cameras. What the home ground confirms in the wider datasets is not numbers but checks: a place where a pipeline decision can be tested against the photographs themselves and against the people who took them, and a small instrument whose daily and seasonal distributions can be verified beside the network's before the two are compared (§3.6). The wider datasets differ from the home ground in ways the home ground cannot calibrate: who placed each camera and why, how far it sees, how sharp its night images are, and who classified its photographs. Two features of the design are meant to explicate those differences rather than hide them. First, because each ratio is formed within one camera, whatever is constant at that camera (its detection distance, its trigger sensitivity, the share of its frames too blurred to name) enters numerator and denominator alike and largely divides out. What does not divide out is anything that changes at a camera over time, or that affects one species' photographs differently from another's, and §§3.5, 3.13 and 3.14 are about exactly those. Second, each claim was tested at the scale that raised it before it was repeated at another (§3.7), and results are reported as agreeing or disagreeing across scales rather than transferred between them; where a result holds only at particular cameras, it is reported as belonging to them (§3.8).

Photographs that could not be named at the class-level - deer frames that show a deer but not its sex - are treated in §3.5.

Interpretation from the ground was aided by the fact that the home-ground numbers were read by a person who also watches the ground they describe (Vraniak 2026g). Twenty years on a second-story deck overlooking the prairie and savanna, often with a long telephoto lens, show turkeys and deer interweaving along the same trails with neither following the other; now and then a

young deer puts its nose near a turkey, which flutters and lands a few feet away. When the statewide numbers later placed the turkey at chance as a follower of deer (§3.8), they confirmed what had been seen. Viewing each of the more than 11,000 home photographs as they were uploaded week by week over three years, and labeling the family's own Snapshot Wisconsin camera on the network's platform, gives some direct sense of what volunteers elsewhere are looking at when they classify. The same viewing supplies context the numbers cannot: two wolf kills in twenty years on the 100 ha (one fawn, one doe); two years of single and paired wolves failing, at night, to take one mature buck at the same junction of three trails; coyotes carrying a fawn and a fawn's leg, a bobcat taking three squirrels in one evening and bobcats carrying turkeys; a bear shifting to the deer's new crossing; the separate loops of doe-and-fawn groups and the different routes of bucks outside the rut; the near-exclusive night use of coyotes and wolves except on rainy days, which bobcats favor; the stillness of most prey on very windy days; the wide berth every species gives a bear; and the time each species spends smelling who has passed. Some fifty years of tracking in the sand of the Northwest Sands Ecological Landscape, in the Lipan Apache tradition as taught by Tom Brown, Jr., adds a further complementary check. Most directly, the person who drives the ATV, the gas mower and the electric cart past each camera, and whose family passes each camera and then sees what happened before and after, knows several of the confounds while the data are being processed. That is how the brushing operation of §3.5 was found, for instance.

3. The issues, in the order the data raised them

3.1 A return interval needs a reference: the species ’own usual interval at the same camera

Problem. A raw time-to-return after a person inappropriately mixes two things: how much the animal may have been affected by the person’s transit, and how often the animal comes to that camera at all. A rare species exhibits a long raw return whether or not it minds people, and a species that visits hourly cannot exhibit a long one. Read raw, the measure ranks species by rarity and reads the ranking as behavior.

What was done. For each species the median gap between consecutive independent detections at the same camera was taken as that species’ baseline at that camera, capped at seven days to keep a single long absence from determining the scale. The baseline is computed afresh at every camera at every scale; no value measured on the home ground is applied to a county or statewide camera. The response index is the median return time after a trigger divided by that baseline. An index of 1.0 means the animal came back on its own schedule; the baseline is reported with the index, so a reader can see what the division was by. A ratio of medians summarizes two distributions, and a broad, flat baseline distribution would make it a poor description. The spread of per-camera ratios is therefore reported with every median (§3.10), and the paragraph that follows sets the delay after a passage against what the camera’s own clock produces when nothing has passed.

An index of 1.0 is a description, not a null. The numerator is a wait measured from an instant the animal did not choose — the moment something else passed — to its next photograph; the denominator is a gap measured between its own visits. Those are different quantities, and nothing makes their ratio 1.0 when the passage has no effect: where the ratio settles depends on

how a species' visits are spaced and clustered at a camera, and it has to be computed. It was computed here on the statewide dataset from the question itself — what should one expect of the next visit if the prior passage had had no effect — in two forms: instants drawn uniformly over each camera's active days, and each passer photograph moved to random active days of its own camera with its time of day kept, which is the construction of §3.8's permutation null. Under the hour-keeping null the deer's expected index is 0.86× after a wolf (383 cameras), 0.84× after a person (4,317) and 0.86× after a turkey (3,499) — below 1.0, so measuring the deer against 1.0 understated its response — while the wolf's expected index is 2.31× after a person (278) and 2.28× after a turkey (234), because the wolf's visits are sparse and clustered and a random instant sits, on the median, more than two of its own visit-gaps from its next photograph. Read against that null, camera by camera, the deer's delay after a wolf is 2.44× what the null expects ($P = 1.8 \times 10^{-41}$), after a person 2.11× ($P < 10^{-300}$) and after a turkey 1.69× ($P = 1.3 \times 10^{-282}$); the wolf's delay after a person is 1.02× ($P = 0.222$) and after a turkey 1.01× ($P = 0.405$). The wolf's 2.42× after a person, carried in the deposits as a raw index, is therefore the null expectation and not a response — which its turkey-referenced excess of 1.06 (§3.16) already said, and which no reader of the raw figure could have known. The turkey reference of §3.11 was then read under the same null: the deer's excess after a wolf over its excess after a turkey is 1.324 observed and 1.019 under the null on the same 315 cameras ($P = 5.3 \times 10^{-5}$ and $P = 0.438$); after a person, 1.169 against 0.987 (3,206 cameras). The structure that moves the raw index away from 1.0 is in both terms of the referenced index and cancels from the difference, which is why the deposited papers read every delay against the turkey rather than against 1.0.

What it changed. On the home ground the red fox's long raw return became rarity once divided, the deer was 0.99 pooled across 2024–2026, and the wild turkey at 1.89 (Vraniak, 2026f). The

same instrument, used at the next level of extract, measured deer at 1.42× in Washburn County, 1.83× across the five counties, and 2.24× after a wolf (Vraniak 2026b) — numbers that are comparable to each other precisely because each was divided by a baseline measured at the same camera.

The lesson is that the count only means something compared with what is normal for that animal at that camera.

3.2 The unit of independence: the pass, not the station, and not the photograph

Problem. A machine pass is not a station-level event. A mowing circuit or a camera check sweeps every station on one trip, and the engine is heard across the whole system at entry. Keying returns to the physical station shatters one whole-system event into per-station fragments and manufactures between-station heterogeneity that has no field counterpart; keying them to the detection pseudoreplicates outright (Hurlbert 1984).

What was done. Human photographs were collapsed into passes, a gap of more than about sixty minutes starting a new pass. Each pass contributes one number: the time to the next deer detection anywhere in the system, divided by the deer baseline. Intervals are bootstrapped on the pass counts, so the reported uncertainty is uncertainty about passes.

What it changed. Station-keyed, the 2025 conveyance contrast fell to $P = 0.059$ and would have been reported as a trend. Pass-keyed it returned clean: on foot 1.00 (71 passes), electric cart 0.95 (65), gas mower 1.12 (15), gas ATV 1.57 (28), with electric against ATV at $P = 0.006$. The earlier per-detection counts (223 against 95) were shown to be pseudoreplicated. The pass unit also made a second pattern legible: exposure was flat across the year (ATV passes, spring ≈12, summer–autumn ≈16) while the response to the same machine roughly doubled from the first

half of the year (≈1.1) to the second (≈2.1). The response is measured by the animal's season, not the machine's schedule. The rule and the share of photographs it keeps at 15, 30 and 60 minutes are in Supporting Information S1.1.The lesson is that the unit of independence is decided by the field, not by the spreadsheet.

3.3 Camera model as the station proxy, and what strict counting does to a P-value

Problem. Where camera serial numbers were not recorded on every frame, the camera model had to stand in for the physical station in the independence rule. Where several cameras share a model, that proxy pools two or more places under one key, and the pooling is invisible in the output.

What was done. The deposited machine paper reported the result under both counting rules and said which rule it used, rather than choosing the rule that gave the smaller number (Vraniak 2026a).

What it changed. Under strict independent-pass counting the pooled electric-versus-gas contrast was $P = 0.14$, and an earlier two-year $P = 0.03$ became $P = 0.18$. The result that survived strict counting was the one keyed to the character of the sound (§3.5), not to fuel. The proxy is priced against habitat in Supporting Information S1.2.

The lesson is to report the counting rule with the P-value; a result that survives only one rule is a lead, not a finding.

3.4 Which instrument each species can honestly be read with: delay for prey, diel overlap for nocturnal carnivores

Problem. A return-time ratio after a person presupposes that the animal and the person could have met. For a nocturnal carnivore whose detections barely coincide with daytime people, the ratio measures the clock and reports it as fright.

What was done. Each species' response index was paired with its coefficient of overlap with people — the intersection of the two hourly activity histograms, after Ridout and Linkie (2009) — and neither was reported without the other. Where overlap was low, the species was read by temporal separation rather than by delay. This was done on the home ground and again in the county extract.

What it changed. On the home ground the raccoon (index 1.63 at overlap 0.20) and the coyote (1.27 at 0.27) looked wary by the index and were interpreted instead as temporally separated, while the turkey (1.89 at overlap 0.74) became the one genuine avoider — an animal with every chance to meet people and a delay that showed it (Vraniak, 2026f). In Washburn County the carnivores' overlap with people fell with nocturnality (wolf 0.17, coyote 0.28, against turkey 0.76), and the deer avoided in both: by delay, and by a nocturnal shift, the county deer ratio rising from 0.70 in spring to 2.32 in the fall firearm season.

The lesson is that metric validity is a finding: know which instrument each species can be read with before reading it.

3.5 Class definitions can bury a signal: fuel against the continuity of a sound

Problem. The first machine classes were measured by fuel, which folded a tolerated gas mower in with a fled gas ATV and averaged the two into nothing. A class boundary drawn on the most available attribute is a hypothesis about mechanism not actually tested.

What was done. The fine machine labels were recovered from the frame table and the classes re-cut by the character of the sound: continuous (electric cart, gas mower) against intermittent (gas ATV). The re-determined boundary was tested before the effect was tested again.

What it changed. Pooled deer return ratios became foot 1.02, electric cart 0.83, gas mower 0.73, gas ATV 1.64, with continuous against ATV at $P = 0.0022$ and does alone at 0.85 against 1.66 ($P = 0.0005$) — where the same dataset was grouped by fuel it only produced $P = 0.14$ (Vraniak 2026a). A scoping pass found the discrimination to be the deer's alone: turkey, coyote, bear and bobcat, all with ample daytime data, showed none of it.

The lesson is that the class boundary is a hypothesis about mechanism, so the boundary is tested before the effect is.

In spring 2026 the frame label "ATV (gas)" covered both an ordinary pass and an extensive brush-cutting operation run from the same machine; no such work was done in 2024 or 2025, so the dataset holds an unplanned contrast. Deer returned at 1.70× baseline after the ATV in March–April 2026 ($n = 16$) compared to 1.23× after the same machine merely passing in the same window of 2025 ($n = 14$), and the response decayed with days since the last ATV day (2.14× within two days, 1.54× at three to seven, 1.30× at eight to twenty-one) — evidence that the delay tracked the operation rather than the altered ground. Identified-doe detections at those

cameras ran 1.10 per camera-day against 1.37 in 2025, so use was reduced while return was delayed, but the animals did not leave.

The lesson is to carry the operation log as a variable: a machine's label names the machine, not what it was doing; the frame labels cannot tell a pass from a day's work, and a study that codes conveyance without coding operation has averaged two stimuli under one name (Vraniak, 2026f).

A class-resolved index is only as stable as the identifications it requires, and the share of frames that can be classified is itself a variable. In this dataset the share of deer frames assignable to sex ran from 100% (April 2025) to 43% (April 2026), the low month falling in the brushing window — clearing a trail edge changes what the camera at that edge can resolve, so the stewardship being measured moved the instrument. On the first pass the doe index after the ATV in March–April 2026 stood at 2.90× compared to an all-deer index of 1.26×, and in five of fifteen ATV events the first deer to return was an unreadable frame, four of them within fifteen hours — quick returns that a sex-resolved index excludes by construction. All 91 unassigned frames were then re-read: 80 were does, three were bucks, two were not deer at all, and six remain unresolved, moving the results for the month from 43% to 96% classified. The doe index fell from 2.90 to 1.57 and the all-deer index rose from 1.26 to 1.70. The excluded frames had been overwhelmingly does, so their exclusion had inflated the doe-specific delay by very nearly a factor of two and had produced an apparent sex effect where the resolved data show a deer effect. The lesson is to report the classifiable share beside every class-resolved index, and to treat a low share as a correctable defect rather than a caveat: where the share is low enough to produce a significant result, re-read the frames and report what the re-reading moved. The direction of the bias is predictable — the frames hardest to classify are not a random sample of the animals in

them — and only the re-reading fixes its size (Vraniak, 2026f; and see Vraniak 2026e, which reaches the same discipline by design rather than by repair).

3.6 Representativeness before widening: the home cameras against the statewide dashboard

Problem. A single site can be idiosyncratic in ways its owner cannot see. Before its numbers are compared to a regional dataset, the small instrument has to be shown to measure the same signal the large one measures — otherwise a difference between the two is a difference between two instruments.

What was done. The Snapshot Wisconsin public dashboard's statewide by-hour and by-month distributions for the well-sampled species were compared to the home dataset's, species by species, before any comparison was drawn between the two records.

What it changed. The home cameras reproduced the statewide diel signature closely — deer overlap 0.91 ($r = 0.94$), coyote 0.90, turkey 0.86 — with two interpretable departures: the home deer are slightly more crepuscular and slightly more autumn-weighted, which is where a quiet property should sit on the published disturbance gradient (Gaynor et al. 2018; Gilbert et al. 2022; Clare et al. 2023).

The lesson is to calibrate the small instrument against the large dataset before the small dataset is allowed to speak about the large one.

3.7 Verification by exact reproduction before any new question

Problem. Each wider extract arrived as a different pull, with its own joins, filters and covariates. A new finding on a new pull can be a property of the pull, and there is no way to tell from the finding itself.

What was done. Before any statewide analysis, the full pipeline was rebuilt from the raw five-county files and required to reproduce every published value exactly; the statewide file was then checked record by record against the five-county file it contains.

What it changed. It did reproduce: 116 sequential encounters, follow fraction 0.75, median lag 5.8 minutes, before/after fractions 0.55 (n = 40) and 0.85 (n = 76), and return ratios of 2.24× (110 cameras), 1.83× (432) and 2.78× (89). The same check caught the one flaw in the delivery: the human-modification covariate had been built from a 90-m snapshot rather than the 1-km native layer and correlated with its parent layer at only r = 0.44. Re-pulled at 1 km it correlated at r = 0.97 with identical medians (0.203), and only then was it used (Kennedy et al. 2019; Vraniak 2026d).

The lesson is to reproduce the last dataset to the digit before the new dataset is allowed to say anything new.

3.8 A directional test needs a null that keeps the clock, and a negative control

Problem. “The wolf came after the deer” can arise from abundance and shared hours alone. Two species that are both common and both nocturnal will appear in each other’s wake at a rate that has nothing to do with pursuit, and a null of 50% assumed rather than constructed cannot separate the two.

What was done. For each focal detection the nearest deer within thirty minutes was scored as co-occurrence, follow or precede, and the follow fraction was compared against a day-reassignment permutation null: 500 iterations, each detection moved to a random day on which its own camera was active in the same season, keeping its time of day. The null therefore holds camera, season, hour and deer availability fixed and varies only the day. The wild turkey, which poses no danger to deer, was then run through the identical test as a negative control.

What it changed. Wolves arrived after deer in 75% of 116 five-county encounters against a null of 50% ($P = 0.002$; Vraniak 2026b) and in 74% of 461 statewide encounters ($P = 0.002$), while the turkey landed at chance — 0.49 against 0.50 across 30,323 turkey–deer encounters ($P = 0.75$) — so the measure reports no following where none exists (Vraniak 2026d). A camera-clustered check on the county dataset then set the honest bound: the directional signals are site-concentrated, while co-occurrence, intervals, class composition and diel structure survive clustering. Window width (15, 30, 60 min) and the permutation seed are examined in Supporting Information S1.3.

The lesson is that a directional result is worth what its null is worth, so a species that cannot be following is carried as the control.

3.9 Confounds tested rather than hedged: shared nocturnality, antler cast, predator density

Problem. Three objections a reviewer would raise against the buck-selection and harvest results could either be hedged in a limitations paragraph or be tested.

What was done, and what it changed. (a) Shared nocturnality. The selection ratio was recomputed within the night alone: bucks were 18.7% of the deer that appeared but 31.2% of those nearest a wolf (1.67, $P = 0.037$), and by day 10.7% against 23.4% (2.20, $P = 0.015$).

Shared hours raise how often a buck is available beside a wolf; they do not create the nearness (Vraniak 2026b). (b) Antler cast. The selection was re-run in the hard-antler months, when an antlered deer is reliably a buck: September–December 1.91 across five counties (Vraniak 2026b), 1.40× statewide ($P = 0.004$) and 1.48× in the October–November rut, against 1.10× — nothing — when the same test is run across the whole year (Vraniak 2026d,e). (c) Predator density. The post-harvest rise in following survived a joint model holding development and latitude fixed (odds ratio 4.8, $P = 0.001$) and a control that subsampled post-harvest wolves down to the pre-harvest count of 559 detections (0.85, band 0.75–0.94, no draw reaching 0.55) (Vraniak 2026c); statewide, the deer's delay after a wolf did not change across a twelve-fold range of local wolf density ($P = 0.22$; $r = -0.005$) (Vraniak 2026d).

The lesson is to test the soft spot; here the exposed result was the one that came back stronger.

3.10 The raw clock against the per-camera ratio

Problem. Pooled across cameras, deer appeared to return sooner after a wolf than after a person, because wolves are recorded where deer are thick. And a researcher who divides the pooled medians (29.9 h ÷ 18.1 h = 1.65×) arrives at a different number from the one reported.

What was done. The ratio was computed inside each camera against that camera's own baseline — a camera qualifying only with at least three baseline intervals and three post-disturber returns, on cameras with at least four deer and two disturber detections — and the reported value is the median of the per-camera ratios. The absolute-time distribution is reported beside it for the reader who seeks the amount of return time (hours).

What it changed. 2.24× after a wolf on 110 cameras (interquartile 1.14–3.20 in the five-county dataset) against 1.83× after a person on 432; in absolute time, a median of 29.9 hours to

reappear, 45% back within 24 hours, 61% within 48, and a long tail reaching about sixteen days at the ninetieth percentile (Vraniak 2026b).

The lesson is to normalize inside the camera and then summarize across cameras, giving the reader the hours as a translation and never as the primary result.

3.11 A floor under every passage-keyed delay: the turkey reference and class-specific floors

Problem. Any return-after-passage measure carries a component that has nothing to do with danger: the plain arithmetic of one animal passing before another. Because that component is larger for animals that appear in bursts of detections, the burstiest deer class can appear to be the most afraid.

What was done. The deer's delay after a turkey — an animal that poses no danger to deer — was taken as the floor, and class-resolved delays were compared against that class's own delay after a turkey, paired within camera and within June–November so that season could not carry the contrast.

What it changed. Deer returned 1.53× more slowly even after a turkey (3,499 cameras). The wolf's 2.19× stands well above that floor and the pooled bear's 1.56× does not, which makes the danger-specific part of the wolf's effect roughly four times the bear's effect. Measured instead against class-specific floors, the bear's delay stood above the reference for every deer class (Vraniak 2026d, 2026e).

The lesson is to price the floor with a species that cannot be feared, and to read every delay as its distance above that floor.

3.12 Designing the before/after: the clean-window comparison for a pulse

Problem. A simple split at the harvest date (before 0.55, after 0.85 in the five counties; Vraniak 2026c) mixes the contaminated harvest year into both sides, and it cannot distinguish a step that persists from a pulse that decays — which is the whole question a management event raises.

What was done. The harvest year 2021 was set aside, and the two clean years before it (2019–20) were compared with the two after (2022–23) and then with 2024–25. The same windows were run for coyote, bear and turkey as a specificity check — species that should not have changed.

What it changed. Following rose from 0.63 to 0.83 (Fisher's exact $P = 0.0019$; odds 2.77× with zone held constant) and returned to 0.68 by 2024–25 ($P = 0.43$): a two-to-three-year pulse, localized to the Northern Forest (0.56 to 0.80, $P = 0.004$). The coyote's fractions in the same windows were 0.61 and 0.62 across 13,461 encounters, and bear and turkey were unchanged (Treves and Louchouarn 2022 for the harvest itself; Vraniak 2026d).

The lesson is that a before/after needs clean windows on both sides, and a species that should not have changed.

3.13 Sex and class must be read in season — and the window must be stated

Problem. Volunteers' deer classes are measurable consistently only in season: antlered means buck reliably in the hard-antler months and unreliably outside them, and a spotless six-month-old is commonly recorded as antlerless. On the home ground, sex came from free-text labels of varying precision. And a result announced as a spring finding can move with the window that defines spring.

What was done. Hard-antler months (September–December) were used for buck identity and spotted months (May–August) for fawns (Vraniak 2026d, 2026e). On the home ground, buck was assigned only where the label implies antlers or a male, doe only where it implies an antlerless female, and mixed frames were excluded. The fawning-window test was run on May–June with March–June carried as a sensitivity check, and the neighboring window was reported whether or not it agreed.

What it changed. The 2026 doe signal — does delaying after people at 1.70 over March–June against bucks at 1.48 — was carried by March–April (1.68, $n = 52$) and June (1.46, $n = 36$) and was absent in May (0.92). On the May–June window alone it was 1.14, about what 2025 gave (1.16). A result that had been measured as fawning vulnerability sits outside the fawning peak (Vraniak, 2026f).

The lesson is to state the window, and to show the result on the neighboring window.

3.14 Labeling schemes across years

Problem. Two annual labeling schemes differed in the rarer categories and in the canid labels. Where a scheme changes between years, a between-year change in the animals and a between-year change in the coding are the same number.

What was done. One harmonization crosswalk was built and applied across years. The ambiguous “coyote or gray wolf” class was carried as coyote in totals and kept separate in every canid analysis, and the 2024 canid records were labeled conservatively (141 ambiguous frames against 21 coyote).

What it changed. The 2024–2025 mesocarnivore rise (6.6% to 11.1% of individuals) fell back to 4.9% in spring 2026 and rests, for 2024, largely on ambiguous frames: a one-year coyote pulse bounded by the labeling rather than measured through it (Vraniak, 2026f).

The lesson is to carry the labeling scheme as a variable.

3.15 Cameras are not independent, and clustering is not a spatial model

Problem. A statewide network of camera data is not a set of independent observations. Volunteers place cameras on the land they have; Snapshot Wisconsin gives each location to the square mile and no closer, so 5,395 cameras occupy 3,232 distinct sections, and neighbors in the same country share weather, deer, roads, ownership and often the same animals. Any covariate measured over an area larger than a camera — and every landscape covariate is — is shared outright between neighbors. Standard errors computed as though the rows were independent are then too small by an amount nobody can guess in advance.

What was done. Standard errors were clustered by survey township — thirty-six sections, deliberately coarser than the nine square miles over which the covariates are summed — and every analysis was repeated with the section as the unit, pooling cameras that share one, and reported beside the camera-level result. The residuals were then tested for the structure clustering does not model: Moran's I under three weightings (cameras within 10 km, inverse distance within 10 km, within 5 km), a distance correlogram in bands, a refit adding Moran eigenvector maps as spatial filters, and, as the crudest and least assuming check, refits of thinned samples holding one camera per section and one per township.

What it changed. It changed a P-value from 0.0009 to 0.59. In the land-tenure analysis (D. A. Vraniak, unpublished data) the statewide wolf rate falls with the public-land share of a camera's

surroundings at a rate ratio of 0.968 per ten points. Computed as though the 2,449 cameras were independent, that estimate carries an interval of 0.950 to 0.987 and $P = 0.0009$. Clustered by section it is 0.864 to 1.084; by township, 0.862 to 1.088 with $P = 0.59$ — nothing. The same correction widens the interval on large private woods from 0.978–1.063 to 0.800–1.300, a fivefold widening on the log scale, and section-level and township-level clustering agree closely, which is itself informative: the non-independence is local. The residual spatial structure is reported in Supporting Information S1.4.

The lesson is to cluster at the scale of the covariate's own support rather than at the camera, to report the unclustered estimate beside the clustered one so a reader can see the size of the correction, and then to test the residuals for what clustering does not model; a correction that turns $P = 0.0009$ into $P = 0.59$ is not a technicality.

3.16 What an intervening species does to an event, and the one thing the threshold cannot touch

Problem. Are two photographs of a wolf, twenty minutes apart, with a photograph of a deer between them, one wolf event or two? If another animal ends the event, the number of wolf events at a camera comes to depend on how busy the camera is — on the deer — and the wolf's rate becomes partly a deer measure. If it does not, a long bout of detections at a busy camera collapses to the same single event as a brief one at a quiet camera. Either way the choice is made once, in the script, and it is almost never reported.

What was done. The rule here is keyed within species. An intervening photograph of any other animal, a person included, neither ends nor restarts an event: two wolf photographs twenty minutes apart are one wolf event whether or not a deer walked through between them. The

reason is the confound above, and the interpretation behind it is that what the intervening photograph records is something about the place, not about whether the wolf left and came back. Where the ordering across species is the measurement rather than a nuisance — the passage-keyed delays of §§3.1 and 3.11, and the sequential encounters of §3.8 — the crossing species is handled explicitly and differently: every passer photograph is counted, none collapsed, and the return is to the next raw photograph of the responder, so the independence rule governs only the responder's own baseline. To estimate the cost these choices the statewide measures were recomputed at fifteen, thirty and sixty minutes, and again with the passer's photographs collapsed at the same threshold as the responder's.

What it changed. It changes the raw index by about a tenth for each doubling of the threshold, and it leaves the referenced index exactly where it is. Statewide, the deer's delay after a wolf is 2.41×, 2.19× and 1.95× the deer's own baseline at fifteen, thirty and sixty minutes (383 cameras at every threshold); after a person, 2.11×, 1.86× and 1.65× (4,317 cameras); after a turkey, 1.74×, 1.53× and 1.35×. The wolf's delay after a person is 2.48×, 2.42× and 2.24×. The movement has one source and it is arithmetic: the threshold enters the index only through the denominator, the responder's own median interval between independent detections, which lengthens as the threshold widens. Collapsing the passer's photographs and changing the anchor of the interval are evaluated in Supporting Information S1.5.

That is also why it cancels. Read as the deposited papers measured it — as the excess over the same animals' delay after a wild turkey at the same camera, paired within camera (§3.11) — the denominator is common to both terms and drops out of the difference. The deer's excess over its turkey reference after a wolf is 1.324 at fifteen minutes, 1.324 at thirty and 1.324 at sixty, on the same 315 cameras, with $P = 5.3 \times 10^{-5}$ throughout; after a person it is 1.169 at all three

thresholds on 3,206 cameras ($P = 3.6 \times 10^{-17}$); the wolf's excess after a person is 1.06 and not significant at any of them. The independence threshold is a free parameter of the raw index and is not a parameter of the referenced index at all.

The lesson is to say what an event is keyed to — the species, the last kept photograph, and whether another species ends it — and to prefer the measure from which that keying cancels.

3.17 The land around a camera: its support, its scale, and whether the cameras can see it

Problem. Every camera study that says anything about landscape joins its cameras to a layer — cover, roads, an index of human footprint, in this case ownership — at some radius, and the radius is chosen by whoever determines the joining of elements. Three assumptions ride on it and none is usually stated: that the layer is measured over an area the analysis is entitled to use; that the area corresponds to something the animal uses; and that the cameras are placed where they can see the class of land being tested. The third is the one is what no amount of care defining the first two actually repairs.

What was done. Snapshot Wisconsin gives a camera's place to the square mile, so the section is the finest honest unit, and exposure was summed over the camera's own section and the eight touching it — nine square miles, a median of 5,706 parceled acres — with shares taken of land parcels rather than of total area, so that lakes and roads do not dilute them. The mismatch between that block and a wolf's territory, which is many times larger, was stated with its direction: the block understates exposure, so a true tenure effect is attenuated toward the null, which runs against a null result rather than for it. The construction of the holdings was then varied — tracts assembled from the parcels of one owner within 25 m, with a strict exact-name variant as a check — and so were the radius of the neighborhood and the definition of the range.

What it changed. The rate ratio for ten points of large private woods is 1.02 as built, 1.04 with tracts assembled from exact names only, 1.06 using the camera's own section alone, 1.11 pooling woods tracts of 160 acres and more, and 1.02, 1.02 and 1.01 with the wolf range drawn at 5, 10 and 20 km. What the checks cannot repair is the third assumption. Of 2,449 cameras, 505 have any large private woods within their nine square miles, 137 have a tenth of the block and 33 a quarter; and because volunteers do not host cameras inside closed private ground, those cameras see the tracts from outside. The design has selection where it most needs sampling, and no radius fixes that.

The lesson is to state the support of every landscape covariate, the scale it is meant to stand for and the direction of the mismatch, and then to ask whether the cameras are placed to see the class of land the analysis names; where they are not, say so in the result and not only in the limitations.

4. A checklist for volunteer-network camera analyses

- Divide every return interval by the responder’s own baseline at that camera, and report the baseline.
- Decide the unit of independence from the field — the pass for a person or machine on a circuit, the thirty-minute visit for an animal — and report the P-value under the strict rule.
- Pair every response index with the responder’s overlap with the disturber; read nocturnal carnivores by overlap, not by delay.
- Test the class boundary — fuel against sound, antlered against buck — before testing the effect, and code what a machine was doing, not only what it was.
- Report the classifiable share beside every class-resolved index, and re-read the frames where that share carries a headline.
- Calibrate a small, verified dataset against the network’s published statewide distributions before comparing the two.
- Reproduce the prior extract’s published values to the digit before asking a new question of a wider one, and check every covariate against its parent layer.
- Use a permutation null that keeps camera, season and hour; run a species that cannot be following as the control; and report the camera-clustered bound.

- Compute ratios inside the camera and summarize across cameras; give hours as a translation.
- Price the floor with a harmless species and read every delay as its distance above that floor, by class.
- For a before/after, set the event year aside, compare clean windows on both sides, and run the species that should not have changed.
- State the seasonal window for any sex- or class-resolved result, and show the neighboring window.
- Carry the labeling scheme as a variable across years.
- Cluster standard errors at the scale of the covariate's own support, not at the camera; report the unclustered estimate beside the clustered one, and test the residuals for the spatial structure clustering does not model.
- Say what an event is keyed to — the species, the last kept photograph, and whether another species ends it — and prefer the measure from which that keying cancels.
- State the support, the intended scale and the direction of mismatch for every landscape covariate, and say whether the cameras can see the class of land being tested.
- Record a permutation null's seed, its number of draws and the pool the draws are taken from.

5. Discussion

What generalizes. None of the seventeen choices is specific to wolves, deer or Wisconsin. Every camera-trap study of avoidance makes each of them, and most make them silently: a unit of independence is chosen when the script is written, a denominator is chosen when the index is defined, a null is chosen when the test is run, and a season is chosen when the classes are read. What the seventeen sections offer is not a new statistic but a reporting discipline — the choice, the alternative, and the number that separates them. Several of the entries move a P-value across the conventional threshold; two of them reverse the reading of a result.

What is local. Three of the findings are this dataset’s own and their generality is a question for other networks: the invariance of the deer’s response across a twelve-fold range of local wolf density, the two-to-three-year pulse in following after the 2021 harvest, and the particular value

of the turkey floor. A second network with a harmless common species and a management pulse of its own could test all three.

Limits. One analyst built every pipeline here, and one region supplied every dataset. The seven-day cap on the baseline is a judgment, defensible but not derived. A return delay measures displacement, not fear: §3.4's two patterns and §3.8's clustered bound mark where that distinction becomes significant. On the home ground the camera model stands in for the physical station, with §3.3's worked example setting out what that costs. And the deer classes in the network data are volunteer classifications, read in season here precisely because they cannot be trusted out of it.

What the photographs hold that the clock does not. Every measure here is a time; none reads the frame. O. Deering and A. Seaborg (University of Wisconsin–Madison, unpublished data) coded the heading of the wolf in the last photograph of each visit and measured it against a landscape feature, and the same coding would turn a sequential encounter from a pair of timestamps into a pair of headings — whether the wolf arriving after a deer faces the way the deer went, and, on the home ground, whether an animal's next visit after a machine moves toward or away from the machine's line of travel. That check belongs to a dataset in which every frame is the family's own rather than to a network extract, and it is set down here as the next thing the home ground can answer.

Data and code availability

Open data and code are not available. The Snapshot Wisconsin extracts analyzed in the deposited papers are held under a Wisconsin Department of Natural Resources data request and are not the author's to redistribute; the home-ground dataset is the family's own. Requests should be

directed to the first author, and to the WI DNR Office of Applied Science for the network extracts.

Acknowledgments

Adrian Treves (UW–Madison Carnivore Coexistence Lab) read the draft of this paper twice, in September and October 2026; his four questions of 20 September became §§3.15–3.17, and his reading of October 2026 is the reason the unit of counting now has one name, the null of §3.1 is computed rather than assumed, and the anchor of the interval is priced in §3.16. He also shared the two camera studies from his lab cited here. The Snapshot Wisconsin extracts were provided by the Wisconsin Department of Natural Resources Office of Applied Science. A large language model (Claude, Anthropic) was used as an analytical tool to re-examine and validate previous statistical analyses; I designed the studies, made every analytical choice reported here, re-read the photographs where that is stated, authored and edited the final text, and take responsibility for its content. Every number reported here was reproduced from the deposited tables or the archived pipeline outputs before anything new was computed.

Ethics statement

No animal was handled, baited or approached for this work; every record is a photograph from a passive, unbaited camera. The home-ground cameras were operated on the author's own land. The Snapshot Wisconsin extracts were provided under a Wisconsin Department of Natural Resources data request. No animal-care protocol or permit was required.

References

Burton, A. C., E. Neilson, D. Moreira, A. Ladle, R. Steenweg, J. T. Fisher, E. Bayne, and S. Boutin. 2015. Wildlife camera trapping: a review and recommendations for linking surveys to ecological processes. Journal of Applied Ecology 52:675–685. https://doi.org/10.1111/1365-2664.12432

Clare, J. D. J., B. Zuckerberg, N. Liu, J. L. Stenglein, T. R. Van Deelen, J. N. Pauli, and P. A. Townsend. 2023. A phenology of fear: investigating scale and seasonality in predator–prey games between wolves and white-tailed deer. Ecology 104:e4019. https://doi.org/10.1002/ecy.4019

Frey, S., J. T. Fisher, A. C. Burton, and J. P. Volpe. 2017. Investigating animal activity patterns and temporal niche partitioning using camera-trap data: challenges and opportunities. Remote Sensing in Ecology and Conservation 3:123–132. https://doi.org/10.1002/rse2.60

Gaynor, K. M., C. E. Hojnowski, N. H. Carter, and J. S. Brashares. 2018. The influence of human disturbance on wildlife nocturnality. Science 360:1232–1235. https://doi.org/10.1126/science.aar7121

Gilbert, N. A., J. L. Stenglein, J. N. Pauli, and B. Zuckerberg. 2022. Human disturbance compresses the spatiotemporal niche. Proceedings of the National Academy of Sciences 119:e2206339119. https://doi.org/10.1073/pnas.2206339119

Hurlbert, S. H. 1984. Pseudoreplication and the design of ecological field experiments. Ecological Monographs 54:187–211. https://doi.org/10.2307/1942661

Iannarilli, F., J. Erb, T. W. Arnold, and J. R. Fieberg. 2021. Evaluating species-specific responses to camera-trap survey designs. Wildlife Biology 2021:wlb.00726. https://doi.org/10.2981/wlb.00726

Karanth, K. U. 1995. Estimating tiger *Panthera tigris* populations from camera-trap data using capture–recapture models. Biological Conservation 71:333–338. https://doi.org/10.1016/0006-3207(94)00057-W

Kays, R., B. S. Arbogast, M. Baker-Whatton, C. Beirne, H. M. Boone, M. Bowler, et al. 2020. An empirical evaluation of camera trap study design: how many, how long and when? Methods in Ecology and Evolution 11:700–713. https://doi.org/10.1111/2041-210X.13370

Kennedy, C. M., J. R. Oakleaf, D. M. Theobald, S. Baruch-Mordo, and J. Kiesecker. 2019. Managing the middle: a shift in conservation priorities based on the global human modification gradient. Global Change Biology 25:811–826. https://doi.org/10.1111/gcb.14549

Lashley, M. A., M. V. Cove, M. C. Chitwood, G. Penido, B. Gardner, C. S. DePerno, and C. E. Moorman. 2018. Estimating wildlife activity curves: comparison of methods and sample size. Scientific Reports 8:4173. https://doi.org/10.1038/s41598-018-22638-6

MacKenzie, D. I., J. D. Nichols, G. B. Lachman, S. Droege, J. A. Royle, and C. A. Langtimm. 2002. Estimating site occupancy rates when detection probabilities are less than one. Ecology 83:2248–2255. https://doi.org/10.1890/0012-9658(2002)083[2248:ESORWD]2.0.CO;2

Naidoo, R., and A. C. Burton. 2020. Relative effects of recreational activities on a temperate terrestrial wildlife assemblage. Conservation Science and Practice 2:e271. https://doi.org/10.1111/csp2.271

O'Brien, T. G., M. F. Kinnaird, and H. T. Wibisono. 2003. Crouching tigers, hidden prey: Sumatran tiger and prey populations in a tropical forest landscape. Animal Conservation 6:131–139. https://doi.org/10.1017/S1367943003003172

Parsons, A. W., C. Bland, T. Forrester, M. C. Baker-Whatton, S. G. Schuttler, W. J. McShea, R. Costello, and R. Kays. 2016. The ecological impact of humans and dogs on wildlife in protected areas in eastern North America. Biological Conservation 203:75–88. https://doi.org/10.1016/j.biocon.2016.09.001

Ridout, M. S., and M. Linkie. 2009. Estimating overlap of daily activity patterns from camera trap data. Journal of Agricultural, Biological, and Environmental Statistics 14:322–337. https://doi.org/10.1198/jabes.2009.08038

Rowcliffe, J. M., R. Kays, B. Kranstauber, C. Carbone, and P. A. Jansen. 2014. Quantifying levels of animal activity using camera trap data. Methods in Ecology and Evolution 5:1170–1179. https://doi.org/10.1111/2041-210X.12278

Sollmann, R. 2018. A gentle introduction to camera-trap data analysis. African Journal of Ecology 56:740–749. https://doi.org/10.1111/aje.12557

Sollmann, R., A. Mohamed, H. Samejima, and A. Wilting. 2013. Risky business or simple solution – relative abundance indices from camera-trapping. Biological Conservation 159:405–412. https://doi.org/10.1016/j.biocon.2012.12.025

Townsend, P. A., J. D. J. Clare, N. Liu, J. L. Stenglein, C. Anhalt-Depies, T. R. Van Deelen, N. A. Gilbert, A. Singh, K. J. Martin, and B. Zuckerberg. 2021. Snapshot Wisconsin: networking community scientists and remote sensing to improve ecological monitoring and management. Ecological Applications 31:e02436. https://doi.org/10.1002/eap.2436

Treves, A., and N. X. Louchouarn. 2022. Uncertainty and precaution in hunting wolves twice in a year. PLoS ONE 17:e0259604. https://doi.org/10.1371/journal.pone.0259604

Treves, A., P. Mwima, A. J. Plumptre, and S. Isoke. 2010. Camera-trapping forest–woodland wildlife of western Uganda reveals how gregariousness biases estimates of relative abundance and distribution. Biological Conservation 143:521–528. https://doi.org/10.1016/j.biocon.2009.11.025

Vraniak, D. A. 2026a. Quiet stewardship and deer return times: sound continuity, not fuel, predicts white-tailed deer avoidance of stewardship machines in a multi-year steward-conducted camera census. EcoEvoRxiv preprint. https://doi.org/10.32942/X2V96R

Vraniak, D. A. 2026b. Patterns of gray wolf and white-tailed deer movement: local, county and regional camera-trap data (2018–2025). EcoEvoRxiv preprint. https://doi.org/10.32942/X2XX09

Vraniak, D. A. 2026c. Wolf–deer movement after the 2021 wolf harvest, by north–south gradient and proximity to towns using a five-county camera record (2018–2025). EcoEvoRxiv preprint. https://doi.org/10.32942/X2D112

Vraniak, D. A. 2026d. Wolf–deer camera-trap detections across Wisconsin's statewide wolf range (2018–2025): a temporary increase in wolves following deer after the 2021 wolf harvest, displacement effects by habitat zone, and avoidance of people across co-occurring species. EcoEvoRxiv preprint. https://doi.org/10.32942/X29X13

Vraniak, D. A. 2026e. Wolf–deer and guild camera-trap detections across Wisconsin's statewide wolf range (2018–2025), resolved by deer class: buck selection in the hard-antler season, bears and fawns, and class-specific reference values for displacement. EcoEvoRxiv preprint. https://doi.org/10.32942/X22H59

Vraniak, D. A. 2026f. Wildlife responses to a resident family: a three-year camera census on restored prairie–savanna (2024–2026). EcoEvoRxiv preprint. https://doi.org/10.32942/X2XQ20

Vraniak, D. A. 2026g. Traditional ecological relationship: family sharing an abundant home within nature requires recurring mutual adaptation rather than the extraction of knowledge by passing visitors. EcoEvoRxiv preprint. https://doi.org/10.32942/X2JT4C